Essays in Corporate Finance and Governance

by

John W. Barry

Business Administration
Duke University

Date:

March 18, 2024

Approved:

John R. Graham, Supervisor

Alon Brav

Simon Gervais

S. Vish Viswanathan

Dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Business Administration
in the Graduate School of Duke University
2024

Abstract

Essays in Corporate Finance and Governance

by

John W. Barry

Business Administration
Duke University

Date:

March 18, 2024

Approved:

John R. Graham, Supervisor

Alon Brav

Simon Gervais

S. Vish Viswanathan

An abstract of a dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in Business Administration in the Graduate School of Duke University
2024

Abstract

This dissertation consists of three essays in corporate finance and corporate governance. In the first chapter, I develop a structural model of executive compensation with non-binding shareholder approval votes ("Say-on-Pay", or SOP). In the model, the Board sets compensation policy and is biased towards offering a high wage (i.e., the CEO influences compensation). Shareholders can fail the vote and punish the Board. The threat of a failed vote affects compensation policy because the Board of Directors internalizes a cost to SOP failure. I show that providing this vote to shareholders has a disciplining effect on compensation and improves firm value. This is not obvious, given that SOP votes rarely fail.

The second chapter presents work with Bruce Carlin, Alan Crane and John Graham. We show theoretically that using a hurdle rate above the cost of capital conveys a bargaining advantage when negotiating with outsiders. This arises through the use of a "sacred" firm-wide hurdle rate for capital budgeting, which allows project managers to commit when negotiating. We give empirical support for our model's predictions.

The third chapter presents work with Murillo Campello, John Graham and Yueran Ma. We study the direct and interactive effects of corporate flexibility (the ability for firms to adapt to changing business conditions) on firm plans at the onset of the COVID-19 crisis. We find that financial, workplace (the ability to easily shift to a remote workplace), and investment (the ability to easily modify investment plans) flexibility played a key role. The key result of the paper is that different forms of flexibility (namely, workplace and investment during the COVID crisis) interact with each other and improve corporate flexibility.

Dedication

Between my finger and my thumb
The squat pen rests.
I'll dig with it.

Seamus Heaney, *Digging*

Contents

List of Tables

List of Figures

Acknowledgements

I would like to thank my dissertation committee Alon Brav, Simon Gervais, John Graham, and Vish Viswanathan. Each has contributed significantly to my dissertation and growth as an academic researcher.

I would particularly like to thank my advisor, John Graham. First, for inviting me to work on his AFA presidential address, which provided an invaluable education on the literature of corporate finance. Second, for giving me a shot and inviting me onto a research project early in my PhD career. Third, for helping me craft my job market paper, from its very early stages to the paper I took and presented on the market. Finally, for the countless hours he has invested in helping me grow as a corporate finance researcher, teacher, and thinker.

I would like to thank Simon Gervais for his guidance and encouragement while crafting my job market paper, and for being endlessly enthusiastic and encouraging.

I would like to thank the co-authors that have helped craft this dissertation: Bruce Carlin, Alan Crane, John Graham, Murillo Campello and Yueran Ma. I have learned a lot from each on how to do research. This dissertation would not exist in its current state without their input.

I would like to thank the friends I have made during my years at Duke: Ali, Levin, Nacho, James, Mory, Greg, Taha, Lin, Sid, Sung, Pei-Chun, Andrew, and so on.

For their love and support, for setting an example on how to be an academic, and for letting me become myself, thanks to my parents.

For lighting up the world and forcing me out of bed every morning, thanks to my dog, Sunny.

Most especially, for being with me from day one, for being a ray of sunshine in my sometimes dreary life, and for everything else and more, thanks to my beautiful wife, Abby.

1

Introduction

This dissertation investigates several topics in corporate finance and governance: shareholder voting and executive compensation, capital budgeting, and corporate flexibility. A common thread across the three chapters is the study of a central component of how firms organize and operate, and particularly if extant governance mechanism and investment practices provide value for shareholders.

In the first chapter, I seek to understand the efficacy of one of the most prominent shareholder voice mechanism in corporate governance. I develop a structural model of executive compensation with non-binding shareholder approval votes ("Say-on-Pay", or SOP). In the model, the Board sets compensation policy and is biased towards offering a high wage (i.e., the CEO influences compensation). Shareholders can fail the Say-on-Pay vote and punish the Board for overpayment. The threat of a failed SOP vote affects CEO compensation policy because the Board of Directors internalizes a cost to SOP failure. However, shareholders themselves internalize a failure cost via dissenting from the Board. I show that providing this vote to shareholders has a disciplining effect on compensation and improves firm value. This is not obvious, given that SOP votes receive over 90% support on average, and rarely fail.

The second chapter presents joint work with Bruce Carlin, Alan Crane and John Graham. We show theoretically that using a hurdle rate above the cost of capital conveys a bargaining advantage when developing a project with outsiders. The advantage arises through the practice of using a "sacred" firm-wide hurdle rate for capital budgeting, which allows project managers to commit to the hurdle rate when negotiating. In empirical analysis, we show that our model's predictions are supported in the data. Our work also has a novel novel empirical result: using a combination of archival and survey data, we show that there is a "shift" in the distribution of realized return on investment (ROI) from just below to just above companies' *ex ante* hurdle rates, which suggests that firms engage in benchmarking of realized returns against the hurdle rate.

The third chapter presents joint work with Murillo Campello, John Graham and Yueran Ma. We study the direct and interactive effects of different dimensions of corporate flexibility on firm plans at the onset of the COVID-19 crisis. Corporate flexibility allows firms to adapt to changing business conditions. Yet, little is understood about the interaction between different aspects of flexibility. The key insight of our work is that, while flexibility is directly important in ameliorating the effects of a crisis, different dimensions of flexibility *interact with each other* to influence firms' plans.

2

Shareholder Voice and Executive Compensation

2.1 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 executive compensation: the Board should set CEO pay to align the interests of management and shareholders, yet directors have an incentive to favor the CEO (Bebchuk and Fried, 2003).

When shareholders disagree with compensation policy and exerting control is not viable, shareholders can convey dissent through *voice* (Hirschman, 1970). 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 compensation decisions and the CEO's performance. As compensation is the primary tool used to limit the agency problem afflicting managerial decision-making, SOP is a potentially important governance mechanism.

Yet the impact of SOP votes is unclear. Compensation policies receive over 90% support on average and only about 7% of SOP votes in the US fail.[1] As Figure 2.1 shows, the

[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 are asked to vote on the CEO's compensation from the

generally positive outcomes of SOP are hard to square with survey evidence in which shareholders express dissatisfaction with CEO pay (Edmans et al., 2021). Likewise, such apparently high SOP support is hard to reconcile with the well-developed literatures studying CEOs' influence on the pay-setting process (e.g., Bertrand and Mullainathan, 2001; Coles et al., 2014) and CEOs' abilities to demand a large share of rents (e.g., Gabaix and Landier, 2008; Cziraki and Jenter, 2022).

An important consideration is that SOP votes are 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 their *ex ante* compensation decision. As Figure 2.1 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 mechanisms by which this influence occurs. In the model, the Board sets CEO pay and is biased towards paying a large wage (i.e., the CEO influences compensation). Shareholders decide to pass or fail the SOP: failure punishes the Board for overpayment, yet may be costly to shareholders.[2] Estimating the model will assess how much boards internalize the cost of SOP failure into its pay decision and whether shareholders consider SOP failure a costly outcome, thus quantifying the influence of SOP on compensation policy.

However, factors beyond these potential costs also influence wage and SOP vote de-

just-passed fiscal year. Throughout the paper, 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 (that is, 30% voting against, see ISS, 2022, Section 5 "Compensation"). 50% is also an important threshold (Hauder, 2019). SOP votes are binding in other countries, such as the United Kingdom. The analysis in this paper of non-binding SOP may not be applicable to those cases, even though the economic forces are likely similar.

[2] SOP votes are *ex post* approval votes on the previous year's CEO compensation, not advisory votes on *proposed* compensation, hence the timing structure of the model (Appendix A.2 and Novick, 2019).

cisions. The size of the Board's overpayment bias is not obvious. Some CEOs are more skilled than others and thus receive higher pay for their effort. The Board and shareholders cannot observe CEO skill directly, and may have different beliefs about the CEO's ability. They learn over time by observing company performance, with each receiving a private signal (like in Taylor, 2010). Quantifying the role 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: 7% in both the simulated and real data. Importantly, it matches the sensitivity of SOP failure likelihood to both the wage and company performance, the primary determinants of SOP vote outcomes (Fisch et al., 2018).

The estimation produces several results. To start, boards are biased towards overpaying CEOs (relative to the profit-maximizing wage), which I refer to as *board capture*. I estimate that the average S&P1500 CEO captures 40.7% of expected surplus, in line with Taylor (2013), who finds that CEOs capture half of the surplus from positive updates about their ability. How does a bias in the pay-setting process 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 cost from SOP failure that is equivalent to 2.06% of firm value.[3] While the unconditional failure rate of 7% means the cost is about 0.14% of value in expectation, the threat of costly SOP failure disciplines CEO pay, even when failure is *ex ante* unlikely. I estimate that SOP as a disciplining mechanism brings CEO wages down by 4.4% on average, in line with Correa and Lel (2016), who find that the adoption of SOP brought

[3] It is important to clarify that these costs 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.

wages down by about 7%. Hence, *shareholder voice affects executive compensation policy.*

Second, failed SOP votes are perceived as costly by shareholders. I estimate that shareholders internalize a cost to SOP failure equivalent to 0.76% of value (about 0.05% in expectation).[4] This aligns with survey evidence from Edmans et al. (2021): shareholders state that failing the SOP may be undesirable, for example because they are hesitant to dissent from the Board on a prominent policy. Though SOP failure is internalized as conditionally costly, my estimates suggest that the disciplining effect of SOP improves firm value by 4.6% on average, consistent with Cuñat et al. (2016), who find that the adoption of SOP increased market value by 5%.

These results highlight the simple economics through which Say-on-Pay votes impact compensation 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. Giving shareholders access to a 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 empirical patterns to reinforce the structural results. As described next, I document several fundamental empirical facts about the Board and shareholder costs to SOP failure which underpin the model.

The first set of new descriptive facts shows that SOP failure leads 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

[4] That is, the cost is equivalent to 0.76% (0.05% in expectation) of each shareholder's equity stake in the firm.

non-fail group). For directors that remain on the board following SOP failure, I also find they are more likely to be removed or step down from the compensation committee: SOP failure is associated with a 1.5 pp increase in the likelihood they are removed from the compensation committee the next year (a 26% larger likelihood than the non-SOP-fail group).

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 such internal and external reputation costs to directors tied to SOP failure.

While directors generally wish to be re-appointed to the Board (the average director salary is around $400 thousand), I argue that a large portion of the Board's perceived SOP 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 issue. Directors 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. From survey evidence in Edmans et al. (2021), shareholders state they avoid SOP failure to maintain relations with the Board; my results above show that SOP failure leads to director turnover. The negative performance evaluation aspect of SOP failure may commit shareholders to raising the rate of Board turnover, which *ex ante* may be undesirable.

Similar motivations exist for the CEO. The model suggests that the SOP is about more than pay, it is a public signal revealing shareholder beliefs about the CEO's match with the firm; the model indirectly predicts that CEO turnover likelihood should be higher when SOPs fail. I find that the turnover rate is around 30% higher in SOP failure relative

to pass (about 9% *vs.* 12%). Given costs associated with CEO turnover (Taylor, 2010), and increased uncertainty for the company and stock price (Clayton et al., 2005), this suggests a motivation for why SOPs rarely fail: shareholders prefer to avoid the negative outcomes associated with CEO turnover.

As further evidence consistent with this cost, I show there is bunching in SOP vote outcomes: defining SOP *dis*approval as one minus the proportion of shareholder that approve the SOP, I uncover excess density directly below the failure thresholds of 30% and 50%.[5] Bunching helps to identify the parameters which measure Board and shareholder costs from SOP failure. 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.

The model recognizes the same force, but highlights that bunching also contains information about the Board cost. The shareholders' voting decision trades off the cost from increasing the threat (probability) of SOP failure against the benefit of reducing the Board's overpayment bias. Importantly, the threat of failure is a function of the wage, so the Board has an incentive to *bunch* wage choices (across realizations of their private information) if the benefit of decreasing the threat of failure outweighs the cost of paying the CEO a lower wage. Bunching thus identifies both the Board and shareholder costs to SOP failure.[6] Similar empirical evidence has been found in Babenko et al. (2019) in the broader context of management proposals. However, to my knowledge, I am the first to use bunching to identify vote failure cost parameters in a structural model.[7]

[5] In SOP, 30% and 50% of shareholders voting against the SOP are key thresholds (Appendix A.2 and Hauder, 2019).

[6] In Appendix A.3.1, I show in a simple model that one can use bunching to uncover a cost associated with SOP failure. This analysis also shows that, though bunching is a useful feature of the data for identification, a richer structural model (such as that presented in Section 2.3 is necessary to separately the Board and shareholder costs.

[7] This methodology has been used in the public finance literature (Saez, 2010; Chetty et al., 2011), and

The empirical evidence of the Board and shareholder cost provides further insight about the economics at play. SOP votes are non-binding, so they do not impact the Board's compensation policy directly. 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.

Finally, the structural model allows me to go beyond estimating parameters and uncovering their implications: I can construct a counterfactual way of implementing SOP. In my baseline model, the Board and shareholders' private information about the CEO determine wage and voting strategies, so SOP vote outcomes are determined in part by different beliefs about the CEO. A simple change to information-sharing in the model emulates a commonly proffered way to engender communication and align beliefs between the Board and shareholders: granting a focal shareholder an advisory seat on the Board (Kakhbod et al., 2023). As explained in Section 2.6, this change allows shareholders to influence *ex ante* proposed compensation policy, as opposed to approving *ex post*.

In the model, this amounts to the Board and shareholders *sharing* their private signals (beliefs) about the CEO *in advance* of their decisions, as opposed to the wage and vote being determined by these possibly divergent signals. In this counterfactual, the SOP failure rate falls, wages decrease on average (though can increase, see Section 2.6) and firm value increases on average. Importantly, this counterfactual does not involve changing structural parameters, these effects are achieved solely by changing the way information is revealed.

The rest of the paper is organized as follows. I first describe the paper's contribution and context within the literature. Section 2.2 describes the data and presents empirical facts about CEO pay and SOP, which both motivate and discipline the model. Section 2.3 presents the structural model and section 2.4 describes the estimation methodology. Sec-

more recently in corporate finance (Antill, 2021; Ewens et al., 2023; Alvero et al., 2023).

tion 2.5 presents the results of the structural estimation. Section 2.6 introduces and analyzes the counterfactual SOP mechanism. Appendix A.1 contains additional results, and Appendix A.2 provides an institutional summary of SOP. Additional empirical, model and estimation details are in Appendices A.3, A.4 and A.5, respectively.

2.1.1 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 (voice) in influencing the decision-maker (the Board), which is directly related to the voice mechanism in my paper — the Board and shareholder costs to SOP failure determine the effectiveness of SOP as a communication device in disciplining wages. Kakhbod et al. (2023) studies how large, passive investors (for whom exit is not a viable path) can act as a coordinating force among the shareholder base.

My empirical results speak to the literature on how non-binding or non-consequential shareholder voting can influence the Board of Directors. Fos and Tsoutsoura (2014) study how 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.[8]

My paper also contributes to the literature on Say-on-Pay. Several papers study the

[8] 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.

effects of the adoption of SOP (e.g., Cai and Walkling, 2011; Ferri and Maber, 2013; Correa and Lel, 2016; Cuñat et al., 2016). These papers show that the increase in voice brought about by implementing 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 in practice: Boards incorporate the threat of SOP failure into the *ex ante* compensation policy, even though the shareholders are hesitant to fail the SOP, supporting the findings on the adoption of SOP.[9]

My paper relates to how institutional investors impact executive compensation. Mehran (1995) and Hartzell and Starks (2003) show a negative relation between blockholder ownership and the level of CEO pay. SOP was introduced in the US explicitly to increase (large) shareholders' ability to monitor compensation policy. My estimates show that SOP is successful in lowering the level of CEO pay, yet SOP failure is costly to shareholders. Several papers have argued that passive investors, generally the largest blockholders, are ineffective monitors due to their tendency to vote with management (Heath et al., 2022). My results show that dissenting from management carries consequences. 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 when they dissent from the Board. Hence, 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 in-

[9] 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 in the run-up to shareholder meetings, suggesting the market internalizes the value effects of vote outcomes in advance.

centives. Lyman (2023) studies CEO turnover and CEO pay policy jointly in a structural model. A structural literature that studies shareholder voting has emerged; e.g., Blonien et al. (2022) study errors in shareholder voting and Pinnington (2022) considers strategic effects in determining passive funds' management biases. 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 in the sense that there is no explicit consequence, the Board's punishment for a negative evaluation arises through a *career concern* or *reputation* channel (Dewatripont et al., 1999). The shareholder cost resembles a hesitancy for the evaluator to give a negative review, which has been explored in the performance evaluation literature (e.g., Bol, 2011). 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, and Blundell et al. (2020) study how the threat of costly regulatory scrutiny curbs pollution.

2.2 Empirical Analysis of CEO Compensation and Say-on-Pay

2.2.1 Data

For the analysis in Section 2.2.2 and the estimation described in Section 2.4, I use data on SOP vote results (Institutional Shareholder Services), CEO compensation (Execucomp), firm accounting data (Compustat), and stock prices (CRSP). The period is 2011-2020, and the sample is S&P1500 firms as I do not observe detailed executive compensation and voting data beyond this group.

Table 2.1 displays summary statistics for the empirical sample. It displays statistics for firm-level variables, CEO-level variables and outcome of SOP votes. The average vote against is about 9% (i.e., the average support rate is about 91%). Only 6.8% of votes fail

(have more than 30% vote against), and 1.8% receive less than 50% support. Firms are on the larger end (due to the focus on S&P1500 firms), so there is a significant right skew in size and revenues.

The average CEO in the sample receives $855 thousand in salary and $138 thousand in bonus. However, bonus is not a strong feature of the sample, with only 15% receiving a bonus greater than zero. An important thing to note — the model is silent on the differences between salary and performance-based compensation; see Page (2018) for structural analysis of the CEO's contract. CEO tenure is 7 years at the median, which will inform the separation probabilities in the estimation.

2.2.2 Empirical Facts

This section documents key empirical facts about SOP and executive compensation that help motivate and discipline the model. Specifically,

1. SOP failure likelihood is driven primarily by CEO pay and company performance.
2. CEOs exert influence over compensation policy via *board capture*.
3. SOP disapproval leads to costly outcomes for directors.
4. SOP voting behavior is consistent with shareholders facing a cost from SOP failure.

Facts 1 and 2 are largely a summary of empirical results known to the literature, collected and framed within my setting, whereas 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 CEO pay and company performance.*

This introductory fact provides a basis for analysis and will inform how the SOP vote is conducted in the model of Section 2.3: the probability of SOP failure is increasing in the level of CEO's wage and decreasing in company and stock performance. Table 2.2 displays regressions in which the dependent variable is an indicator for SOP failure (more

than 30% voting against the SOP, in columns 1-4), or the percentage of shareholders voting against the SOP (the continuous measure, columns 5-8). The main independent variable is log CEO compensation. The table shows that SOP failure likelihood and SOP disapproval rates are strongly increasing in the level of CEO pay. This relation is robust to performance controls and fixed effects (even as fine as firm × CEO), as well as lagged pay.

The same table shows that SOP disapproval is *decreasing* in company and stock performance (the firm's return on assets (ROA) and the 12-month stock return, respectively), conditional on the CEO's pay, confirming findings from Fisch et al. (2018).[10] These strong relations provide clarification for the quantitative model. Shareholders fail SOPs when wages are (too) high, given what they believe about the CEO. If the company is doing well, then shareholders are less likely to fail the SOP, even if wages are high. Both of these forces will inform the structure of the SOP vote in the model.

As corroborative evidence that SOP impacts *future* compensation policy, Appendix Table A.2 shows a strong negative relation between changes in CEO pay and SOP disapproval: CEO pay falls by about 4 percentage points following SOP failure. This robust result suggests that SOP failure pushes the Board to make changes to CEO pay.

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

CEOs can partially determine the compensation-setting process by influencing the Board of Directors (e.g., Graham et al., 2020). Concurrently, CEOs have bargaining power over the firm (e.g., Taylor, 2013). If primarily the latter, then inflated wages may be optimal, so confirming the presence of the former in my data (which SOP is designed to combat) is important. I examine a well-established measure 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 my data

[10] Table A.1 tests the company/stock performance hypothesis separately.

in Table 2.3 Panel A. The level of CEO pay increases with board co-option: pay increases by 7-9 percentages points for each standard deviation increase in board co-option. As in Coles et al. (2014), I include CEO tenure fixed effects in each specification as co-option mechanically rises with tenure.

Panel B of Table 2.3 presents a result new to the literature. I regress log changes in CEO pay (from t to $t + 1$) on an interaction between board co-option and the outcome of the SOP vote (both from year t). The table shows that board co-option modulates the relation between changes in CEO pay and SOP disapproval. In other words, higher board capture lessens the influence that SOP has on compensation policy.

The model incorporates CEO influence on the Board directly into the Board's pay-setting process. In the model, the Board wants to overpay the CEO by a constant proportion (determined by a parameter λ, which will be discussed in detail in Section 2.3). However, as I show in Appendix A.4.1, neither the shareholder nor the econometrician is able to separately identify whether board capture is CEO influence or CEO bargaining power.; in subsample estimation (Section 2.5.4), I show that estimated board capture varies with board co-option (the measure from Coles et al., 2014).

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 2.4 Panel A shows that director turnover correlates with SOP disapproval. I identify turnover events occurring between SOP votes and regress a director turnover indicator on the SOP vote result from the previous year.[11] Columns 1-3 of Panel A show that director turnover likelihood is 1.5 to 2.3 percentage points (pp) higher after SOP failure. Relative to non-failure, the probability of director turnover is 20% higher. This finding is robust to controlling for company performance (ROA and firm's stock return over the past 12 months), along with a battery of controls

[11] Director turnover is identified following the methodologies in Fischer et al. (2009) and Iliev et al. (2015).

covering board composition, and director and CEO features.

Panel B of Table 2.4 presents a second cost to directors from SOP disapproval. Focusing on the subsample of directors not turned over, I identify cases where compensation committee members leave the compensation committee. SOP failure is associated with a 1.1 to 1.5 pp higher probability of leaving the compensation committee (a 26% increase). Taken together, Panels A and B show that SOP failure is a career concern for directors.

Panel C shows a third cost to directors by evaluating how SOP disapproval affects the *external reputation* of directors — via its effect on the number of outside boards that the director sits on. I focus on compensation committee directors that sit on at least one outside board and then see a reduction in outside board seats the year after the SOP vote. The panel shows that SOP failure is associated with a 1.7 to 1.9 pp increase in the likelihood that the director loses at least one of these outside board positions (a 21% increase). Failing the SOP impacts directors *outside the firm where they work*: external reputation costs result from SOP failure.

Table 2.4 shows that failing SOP votes is a career and reputation concern for compensation committee directors. In the model, SOP impacts CEO pay policy through SOP failure being costly to the Board. The evidence suggests that the cost (labeled χ_B) is large. These findings support Fos and Tsoutsoura (2014) and Aggarwal et al. (2019), who respectively find that proxy contests have lasting reputational damage on directors and abstained votes in uncontested director elections lead to negative consequences for directors.

Fact 4. *SOP voting 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. (2021) provide survey evidence that of the cost: in interviews, institutional investors express their reluctance to fail SOP

votes.[12] 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.[13]

In the model, SOP failure conveys that shareholders believe the CEO is of low skill or the match quality is poor. An implication from the model is that CEO turnover likelihood should positively correlate with SOP disapproval. However, changing the CEO is costly: Taylor (2010) estimates turnover costs equivalent to 1.3% of assets and Clayton et al. (2005) show that CEO turnover leads to long-term increases in stock return volatility. Hence, reducing the threat of SOP failure and ceding some value through higher CEO wages may be preferable.

In support of this, Table 2.5 displays a new result: SOP disapproval is associated with large increases in the likelihood of CEO turnover. When SOP votes fail, CEOs 2.3 to 3.2 pp more likely to be turned over, about a 30% increase relative to turnover rate when SOPs pass. This result is robust to nonlinearities in firm/stock performance and fine fixed effects, suggesting that SOP disapproval signals Shareholder dissatisfaction with the CEO.

While the survey and CEO-turnover based evidence are suggestive of the existence of a cost to shareholders from failing SOPs, the model allows me to use bunching of SOP vote outcomes below failure thresholds to reveal information on this cost. A higher occurrence of close passes relative to fails suggests that shareholders strategically avoid failing SOPs.[14]

[12] It is important to note that Edmans et al. (2021) survey UK institutional investors, so the respondents are not discussing their views on SOP in the US. SOP votes in the US and UK receive similar levels of support.

[13] See Online Appendix A of Edmans et al. (2021). Investors also mention they often follow proxy advisors as resource constraints prevent them from fully analyzing compensation policy. Figure A.1 presents anecdotal evidence of future monitoring costs: after failing the 2021 SOP, Netflix engaged with large shareholders about compensation numerous times throughout the year.

[14] This type of strategic behavior requires coordination across diffuse, or. large pivotal blockholders swing-

Figure 2.2 displays a result new to the literature. I 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. There is clear bunching and I find a statistically significant difference in density.[15] Although bunching is not definitive evidence of a shareholder cost from failed SOP votes, it is consistent with this cost. Pivotal blockholders have an incentive to pass a close vote if the cost of SOP failure outweighs its benefit. Once it is recognized that the outcome of the vote is a function of the Board's wage policy (see Section 2.3.3), bunching also provides information on how costly failed votes are to Boards.

To further motivate the existence of these costs, in Appendix A.3.1, I present a simple structural model (independent of the model in Section 2.3) which uses the observed bunching in Figure 2.2 to estimate a reduced-form object related to the costs of SOP failure. Importantly, this simple model does not place enough structure on the data and *cannot separately identify* the Board and shareholder costs to SOP failure: the object maps to a combination of the two. Nevertheless, the outcome of the estimation in Appendix A.3.1 is similar in magnitude to the costs I estimate in Section 2.4 (both Board and shareholder), which is reassuring.

The identification of the model's SOP failure cost parameters uses the observed bunching. In the model in Section 2.3, shareholders weigh the expected cost of SOP failure against its impact on the Board's wage decision when SOP failure is *ex ante* more probable. In turn, the Board is aware of the distribution of shareholders' beliefs and understands that in some states, even slight reductions in wages can significantly reduce the likelihood of the vote failing. While the underlying signal distributions for both parties

ing the outcome by keeping the percentage of dissenting votes below the failure threshold.

[15] Figure A.3 conducts a test where I set "placebo" vote failure thresholds of 20% and 40%, as opposed to the commonly accepted thresholds of 30% and 50%. The figure shows no change in density at these thresholds.

are continuous, the model predicts that there will be a clustering of vote outcomes directly below the failure threshold with a corresponding gap in the distribution directly above it.

This section presents several new stylized facts about CEO pay and SOP. These facts discipline the model (Section 2.3), and inform the estimation (Sections 2.4-2.5). Particularly important are the motivations for model parameters which embed the Board's bias towards over-paying the CEO (board capture), and the magnitude of the Board and shareholder cost to SOP failure.

2.3 Model

This section outlines the model. The key forces are guided by the analysis in Section 2.2. The subsequent estimation will determine the extent to which these forces matter; for now, the model treats them as parameters. Table A.3 displays model parameters and definitions. The framework, particularly the belief formation process, is inspired by Taylor (2010).

2.3.1 Technology and Environment

Time is annual and the firm is infinitely-lived.[16] The firm consists of three actors: the Board of directors, that sets the CEO's wage 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"). CEO skill is uncertain: the Board and shareholders form beliefs based on the information they observe. Each period, a CEO of tenure τ separates from the firm with (exogenous) probability $f(\tau)$; upon separation, the firm matches (exogenously) with a new CEO.

[16] Though SOPs only need to occur once every three years, in practice most firms have them annually. I abstract away from this — if a firm has an SOP every two year three years, in the data analysis, I take only the year that an SOP occurs (rather than averaging across years), but either method works just as well.

Effort (n_t) is increasing but concave in the wage 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

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

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.[17] Output is thus

$$y_t = A_t w_t^{\alpha}. \tag{2.1}$$

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

$$\ln A_t = a + \varepsilon_{yt}, \quad \varepsilon_{yt} \sim N\left(0, \sigma_y^2\right). \tag{2.2}$$

Type a is not observed by the Board and shareholders, they make predictions about a based on information they observe. Eq. (2.2) defines the notion of CEO skill — higher types achieve higher average productivity. Net operating income in year t is revenue minus the CEO wage:[18]

$$\pi_t = A_t w_t^{\alpha} - w_t.$$

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

[17] Appendix A.4.1 presents a simple microfoundation of the relation between CEO effort, the wage and output. Beyond this, the CEO is a passive actor: the model is silent on the 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.

[18] Profits (by assumption) are paid out immediately as dividends (including negative profits), and thus firm size is fixed over time and will not factor into the model (Taylor, 2010).

SOP, the Board would choose w_t to maximize

$$\pi_t^{\mathrm{B}} = A_t w_t^{\alpha} - w_t + \lambda w_t,$$

where $\lambda \in [0, 1)$ 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 2.2.2 and Fact 2). When $\lambda > 0$, the Board's optimal wage is above that which maximizes net operating income.[19] This bias can reflect a large outside option or some other channel which allows CEOs to demand a high wage (Cziraki and Jenter, 2022; Gabaix and Landier, 2008). It can also reflect CEO influence on the pay-setting process, for example via personal relationships with members of the Board (Graham et al., 2020; Coles et al., 2014). The model does not attempt to separate these forces, and in Appendix A.4.1 I show they are in fact not separately identified; rather λ measures a gap between the wage that would maximize shareholder value and the wage the Board would pay the CEO.

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

$$AC(w_t; w_{t-1}, \tau_t) = c_w \times w_t \left(\frac{w_t - w_{t-1}}{w_t} \right)^2 \times \mathbf{1}[w_t < w_{t-1}] \times \mathbf{1}[\tau_t > 0]. \tag{2.3}$$

The adjustment cost is quadratic, scales with the wage level, is not present if the CEO is in the first year of their tenure, and only activates if the Board decreases the wage from $t-1$ to t; the one-sidedness of the cost is chosen to match the observed downward rigidity in CEO wages. Parameter c_w controls the cost of adjustment, and is to be estimated.[20]

[19] I could write biased operating income as $(1 - \nu)(A_t w_t^{\alpha} - w_t) + \nu w_t$, so the Board maximizes a weighted average of profits (shareholder value) and the extra wage it pays the CEO. It is equivalent to define $\lambda = \frac{\nu}{1-\nu}$.

[20] Table A.2 and the findings from Taylor (2013) show why the adjustment cost is needed. The table shows that CEO pay falls when SOPs fail, whereas Taylor (2013) shows CEO pay is downward rigid; SOP failure *forces* the Board to alter their compensation policy and face the negative effects of lowering the CEO's wages.

Shareholders hold an approval vote each year on the Board's CEO pay policy, or a Say-on-Pay (SOP).[21] The vote is non-binding (it does not force the Board to set a new wage contract) and occurs at the end of each period (after the wage decision and output have occurred). I assume that 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

$$A_t w_t^{\alpha} - w_t + \lambda w_t - \chi_B \times \mathbf{1}[\text{SOP fail}_t] - AC(w_t; w_{t-1}, \tau_t). \tag{2.4}$$

Shareholders in the model seek to maximize net operating income. Thus, if wages are "too high" given shareholders' current beliefs about CEO ability, the SOP vote may fail and the Board must incur the failure cost. Upon vote failure, the shareholders will also face a failure cost. As discussed in Section 2.2.2, such cost might represent an aversion to dissenting from the Board on compensation policy.[22] The parameter $\chi_S \geq 0$ governs the shareholders' cost of failing the SOP. The shareholders' per-period utility is

$$A_t w_t^{\alpha} - w_t - \chi_S \times \mathbf{1}[\text{SOP fail}_t]. \tag{2.5}$$

Eqs. (2.4) and (2.5) summarize the differences in Board and shareholder preferences: the Board is biased towards paying a higher wage, their SOP failure costs may differ, and the Board incorporates the one-sided adjustment cost.[23] Shareholders can alleviate the Board's bias (λ) by threatening to fail the SOP and forcing the Board to pay a (perceived) cost (χ_B); however, they also internalize their own (perceived) cost from failed SOP votes (χ_S). The parameters χ_B and χ_S and their respective magnitudes will be the focus of the estimation.

[21] While the structure of the CEO's pay package certainly influences SOP outcomes, shareholders predominantly vote in response to the level of CEO pay, as detailed in Fact 1 and Table 2.2.

[22] Like the Board cost, in the model, the cost is a perceived aversion to failing the SOP.

[23] The assumption that shareholders do not incorporate the adjustment cost keeps the shareholders' problem static, which simplifies the numerical solution. Moreover, Table 2.2 shows that lagged log pay does not have an impact on SOP outcomes: it is likely that shareholders vote in a "static" sense.

2.3.2 *Model Timeline, Signals, and Beliefs*

This section describes the per-period model timeline (Figure 2.3) and the points at which the Board and shareholders observe information which updates their beliefs about CEO skill a. To start, each period the firm separates from the CEO with exogenous probability $f(\tau)$ and matches with a new CEO of tenure $\tau = 0$. Upon matching, the Board and shareholders begin with a shared prior about the CEO's skill,

$$a \sim N\left(\mu_0, \sigma_0^2\right), \tag{2.6}$$

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,

$$z_{bt} = a + \varepsilon_{bt}, \quad \varepsilon_{bt} \sim N\left(0, \sigma_{z_b}^2\right). \tag{2.7}$$

z_{bt} represents *operational interaction* with the CEO, and it will inform the Board's wage decision. Shareholders do not observe z_{bt}, which means that the Board is asymmetrically informed (has different 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 shareholder in the continuum of shareholders 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. However, in shareholder voting, signals are correlated (this correlation could reflect proxy advisors' recommendations.) This correlated signal structure among the firm's shareholders makes it informationally equivalent (from the econometrician's perspective) to focus on a representative shareholder, as microfounded in Appendix A.4.2. 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:

$$z_{st} = a + \varepsilon_{st}, \quad \varepsilon_{st} \sim N\left(0, \sigma_{z_s}^2\right). \tag{2.8}$$

At the same point, 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 (eq. 2.2), B and S will revise their beliefs about the CEO upward. I label this signal:

$$z_{yt} = \ln A_t = a + \varepsilon_{yt}, \quad \varepsilon_{yt} \sim N\left(0, \sigma_y^2\right). \tag{2.9}$$

Both the private signal z_{st} and the public signal z_{yt} will affect the SOP result, which incorporates company performance into the vote outcome (Fact 1 and Fisch et al., 2018).

Model timeline Figure 2.3 displays the sequence of events within the model. At the start of each period, B and S each believe that CEO ability $a \sim N\left(\mu_{at}, \sigma_{at}^2\right)$. As described in Appendix A.4.4, due to signal disclosure in the wage-SOP game, B and S share the same beliefs at the beginning of each period because B's wage choice fully reveals their private signal and the outcome of the SOP vote fully reveals S' signal.

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 and output realize. 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 they see signals. I use the subscript a to refer to beliefs shared by B and S, so $\left(\mu_{at}, \sigma_{at}^2\right)$ refers to shared beliefs at the beginning of period t. I use the subscripts b and s when B and S can have different beliefs. CEO tenure fully determines the variance of beliefs (see Appendix A.5.2 and Taylor, 2010), so I define the function $\sigma_a^2(\tau)$ to track how the variances of beliefs decreases across tenure. Given the CEO's tenure τ_t in year t, the

variance evolves according to

$$\sigma^2_{at+1} = \sigma^2_a(\tau_t + 1) = \sigma^2_0\left[1 + (\tau_t + 1) \times \sigma^2_0\left(\sigma^{-2}_{z_b} + \sigma^{-2}_{z_s} + \sigma^{-2}_{y}\right)\right]^{-1}. \tag{2.10}$$

The mean evolves according to

$$\mu_{at+1} = \sigma^2_a(\tau_t + 1)\left[\frac{\mu_{at}}{\sigma^2_a(\tau_t)} + \frac{z_{bt}}{\sigma^2_{z_b}} + \frac{z_{st}}{\sigma^2_{z_s}} + \frac{z_{yt}}{\sigma^2_y}\right]. \tag{2.11}$$

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

2.3.3 The Say-On-Pay Vote

Each period, the Board sets the wage and the Shareholder decides whether to fail the SOP. These decisions are informed by each party's private signal and how they view each other's beliefs. This section details S' strategy, holding the B's wage choice fixed. In Appendix A.4.4, I detail several assumptions made about the SOP vote; these assumptions are intended to keep the game tractable, while also remaining realistic about B and S' strategies.

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

$$k_{st} = s_t \times w_t. \tag{2.12}$$

Fact 1: Pr(SOP fail) is increasing in the wage Eq. (2.12) 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 2.3, 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

$$Z_{syt} = pz_{st} + (1-p)z_{yt} = a + p\varepsilon_{st} + (1-p)\varepsilon_{yt}, \quad Z_{syt} \sim N\left(\mu_{at}, \sigma_{at}^2 + \frac{\sigma_{z_s}^2\sigma_y^2}{\sigma_{z_s}^2+\sigma_y^2}\right). \tag{2.13}$$

This is the distribution that matters for 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 (2.13). Further, B knows this is the distribution that from which S' signal will be drawn; B takes this into account when setting the wage.

Fact 1: Pr(SOP fail) is decreasing in company performance Eq. (2.13) is how firm performance affects SOP outcomes in the model: better performance leads to a lower probability of failure (Fisch et al., 2018). In the model, if the firm's 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 is

$$w_t^{\mathrm{U}} = \arg\max_{w_t} \quad E_{st}\left[\exp\left(a + \varepsilon_{yt}\right)\right] w_t^{\alpha} - w_t = \left(\alpha E_{st}[\exp\left(a + \varepsilon_{yt}\right)]\right)^{\frac{1}{1-\alpha}}. \tag{2.14}$$

The goal of the Shareholder in the model is to get the Board's wage choice as close to the unbiased wage as possible, given the costliness of SOP failure. SOP failure should be determined via *distance* in the observed wage from the unbiased wage. The random variable w_t^{U} is lognormally distributed, with distribution determined by belief tuple $\left(\mu, \sigma^2\right)$ and parameters

$$w_t^{\mathrm{U}} \sim \log N\left(\frac{\mu}{1-\alpha} + C, \frac{\sigma^2}{(1-\alpha)^2}\right), \quad C = \frac{\log\alpha + \frac{1}{2}\sigma_y^2}{1-\alpha}. \tag{2.15}$$

Thus, given the distribution of Z_{syt} in (2.13), there is a random variable $w_t^{\mathrm{U}}\left(Z_{syt}\right)$ given by (2.15) that is the conversion of Z_{syt} to its unbiased wage counterpart.[24] I refer to the CDF of this distribution as F_{st}^{U}, where st signifies the Shareholder's period t beliefs about CEO ability.[25]

Determining the probability of SOP 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

$$\Pr(\text{SOP fail}_t) = \Pr\left(w_t^{\mathrm{U}}\left(Z_{syt}\right) \leq s_t \times w_t\right) = F_{st}^{\mathrm{U}}(s_t \times w_t). \tag{2.16}$$

Fixing the Board's best response w_t for now, S chooses s_t to maximize expected net operating income, conditional on their beliefs at the start of period t, $a \sim \left(\mu_{at}, \sigma_{at}^2\right)$. Importantly, S influences expected wages. When setting the vote policy, conditional on the

[24] 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 (2.15).

[25] Conveniently, the transformation to the lognormal, unbiased wage distribution ensures that S' strategy s_t is just a non-negative, real number, and that the threshold is always increasing in the wage. The wage is guaranteed to be positive (due to lognormal productivity), however CEO ability can be both positive and negative (and hence, draws from the distribution of S beliefs about CEO ability can be both positive and negative.

shared belief μ_{at} at the beginning of period t, S takes expectations over signals z_{bt}. At the same time, given z_{bt}, B updates beliefs to $\mu_{bt|z_b}$ following standard updating. B then optimally offers $w_t(z_{bt}, s_t)$, which will be detailed in Section 2.3.4. The Shareholder's problem is

$$\max_{s_t} \quad \int_{z_b} f(z_b) \left[\underbrace{E_{st}[A_t] w_t(z_{bt}, s_t)^{\alpha} - w_t(z_{bt}, s_t)}_{\text{Expected operating income}} - \underbrace{\chi_S \times F_{st}^{\mathrm{U}}(s_t \times w_t(z_{bt}, s_t))}_{\text{Expected cost of SOP failure}} \right] dz_b. \quad (2.17)$$

The Shareholder commits to an *ex ante* threat of vote failure intended to reduce the Board's bias that is increasing in the wage and decreasing in observed productivity. The realizations of z_{bt} (determining the wage), z_{st} and z_{yt} will determine the outcome of the vote.

2.3.4 The Compensation Committee

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

$$\begin{aligned} \mu_{bt|z_b} &= \sigma_{bt|z_b}^2 \left(\frac{\mu_{at}}{\sigma_{at}^2} + \frac{z_{bt}}{\sigma_{z_b}^2} \right) \\ \sigma_{bt|z_b}^2 &= \left(\sigma_{at}^{-2} + \sigma_{z_b}^{-2} \right)^{-1} = \frac{\sigma_{at}^2 \sigma_{z_b}^2}{\sigma_{bt}^2 + \sigma_{z_b}^2}, \end{aligned} \quad (2.18)$$

which is standard Bayesian updating (see Appendix A.4.3). If B revises its beliefs about CEO ability downwards, they will want to decrease w_t relative to w_{t-1}. To match dynamics in wages, I include the wage adjustment cost from (2.3) in the Board's problem. At the compensation committee meeting, B chooses the wage for the upcoming year by solving the following dynamic program

$$\begin{aligned} V(\mu_{at}, \tau_t, w_{t-1}) = \max_{w_t} \quad & E_{bt|zb}[A_t] w_t^{\alpha} - w_t + \lambda w_t - \chi_B \times F_{st}^{\mathrm{U}}(s_t \times w_t) - AC(w_t, w_{t-1}; \tau_t) + \\ & \delta_B \Big[(1 - f(\tau_t)) E_{bt|zb}[V(\mu_{at+1}, \tau_t + 1, w_t)] + f(\tau_t) V^R \Big]. \end{aligned} \quad (2.19)$$

The state consists of the two variables that track the Board's beliefs about the CEO: the current belief about the mean μ_{at} and the CEO's tenure τ_t (which determines the variance of beliefs σ^2_{at}, eq. 2.10). The third is the previous period's wage w_{t-1}, which tracks the history of wages paid to the CEO (which is set to zero if $\tau_t - 1 = 0$). To be clear, w_t is a function of both z_{bt} and the shareholder's strategy s_t, however for tractability this notation is omitted from (2.19). $F^{\mathrm{U}}_{st}(s_t \times w_t)$ specifies the probability of SOP failure as detailed in Section 2.3.3. 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. 2.18).

At the heart of the Board's problem is the trade-off between paying a higher wage (board capture λ) against the higher probability of SOP failure. The Board must calculate the expected productivity of the firm A_t, (also influenced by the distribution of ε_{yt}); the likelihood the Shareholder will fail the SOP vote; the adjustment cost; and the continuation value. The hazard function $f(\tau_t)$ controls how often the firm separates from the CEO and is an input to the model. Pooling all CEO spells, $f(\tau_t)$ is the frequency of turnover after τ years, conditional on the CEO surviving $\tau - 1$ years.[26] V^R describes the termination value. Upon separating from the CEO, the Board accesses the CEO talent pool and beliefs reset to (μ_0, σ_0),

$$V^R = V(\mu_0, \tau = 0, 0), \tag{2.20}$$

so the value function returns to its starting value (prior beliefs) and there is no previous wage.

[26] 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 $\mathbf{T}$ is the cap on the length of CEO tenure. Further, the estimation sample excludes CEO spells that only last the first year.

Objective Firm Value

The Board's optimal wage policy admits the model's definition of objective firm value. Given the state and the Board's choice of the wage, firm value is

$$V^{\text{OBJ}}\Big(w_t, \mu_{at}, \tau_t, w_{t-1}\Big) = \pi_t - AC(w_t; w_{t-1}, \tau_t) + \delta_B \times E_{at}\Big[V^{\text{OBJ}}\Big(w_{t+1}, \mu_{at+1}, \tau_t + 1, w_t\Big)\Big], \tag{2.21}$$

where $\pi_t = E_{at}[A_t]w_t^{\alpha} - w_t$ is the firm's true operating income given w_t. In contrast to the Board's problem, objective value does not include board capture. In general, the Board's policy w_t will be above the value-maximizing wage. Secondly, firm value does not contain any inputs from the SOP vote, it merely represents the discounted cash flows of the firm. When I undertake counterfactuals, (2.21) will allow me to analyze how changes to SOP affect firm value (via its effect on the Board's compensation policy).

2.3.5 Model Solution and Predictions

I use Bayes' rule to derive Board and shareholder beliefs about CEO ability and substitute these beliefs into the Board's Bellman equation (2.19) and the Shareholder's objective function (2.17). Appendix A.4.5 gives the full derivation of the solution. I now discuss the model's predictions relating to the Shareholder's SOP-voting strategy. I focus on the static version of the model, so the adjustment cost and continuation value in (2.19) are dropped and (2.17) remains the same.

The Shareholder's Strategy

To illustrate how the SOP vote influences Board compensation policy, I explore how varying the Shareholder's strategy influences model outcomes. Figure 2.4 displays comparative statics of the shareholder's choice of s. For each point, I set S' strategy to $s \geq 0$, and then plot B's best response *as if* the threshold S plays were the equilibrium.

Panel A shows the Shareholder's maximization problem as a function of their strategy s; the solid blue line (left y-axis) displays Shareholder expected utility. The Share-

holder balances the expected cost of failure against bringing the wage closer to the profit-maximizing wage. The dashed orange line (right y-axis) displays the Board's expected utility. If S' threshold is set to zero, the Board sets the fully biased wage. Shareholders undo the effects of board capture λ via increasing the *ex ante* probability of failure. Panel B displays the expected probability of failure, along with the optimal s from Panel A (the black dashed line, about 10%). Panel C shows the expected wage as a function of the strategy. In this example, S brings wages down by about 15%. Panel D displays operating income $\pi = E[y - w]$, showing the effect of the Shareholder cost: operating income would be closer to its maximum if the cost were smaller.

2.4 Estimation

I estimate the parameters of the structural model presented in Section 2.3 using indirect inference (McFadden, 1989; Smith, 2016): 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 (from the structural model). The auxiliary model, although misspecified, focuses on features of the data which are highly informative about structural parameters. Details of the estimation are presented in Appendix A.5. Section 2.4.1 outlines the identification strategy. I present the quality of the model's fit in Section 2.4.2 and present the estimated structural parameters (and their economic implications) in Section 2.5.

2.4.1 Identification Strategy

I show there is a tight relation between the reduced-form outcomes of the auxiliary model and structural parameters, which is key to the success of the indirect inference approach. While identifying each structural parameter is important, the discussion will focus on the parameters λ (board capture), and χ_B and χ_S (the Board and Shareholder SOP failure costs, respectively). Table A.7 displays the reduced-form outcomes and the parameter(s)

each outcome targets.

Output and CEO skill parameters In Appendix A.5.1, I show how I extract the CEO component of company revenues, which is useful for identification:

$$\log y_{it} = \log \eta + a_i + \alpha \log w_{it} + \varepsilon_{yit} \implies y_{it} = \eta A_{it} w_{it}^{\alpha} = \eta \exp\left(a_i + \varepsilon_{yit}\right) w_{it}^{\alpha}. \tag{2.22}$$

y_{it} is CEO-driven firm revenues and w_{it} is the CEO's observed compensation. η is a factor which scales output to its appropriate level and is important in the estimation as it allows for appropriate comparison across subsamples (Page, 2018). Average CEO skill is only identified relative to the constant $\log \eta$ in (2.22), so I normalize μ_0 to zero.

The following (pooled) regression in the data maps exactly to company output in the model

$$\log y_{it} = y_0 + y_1 \log w_{it} + \epsilon_{it}^{y}, \tag{2.23}$$

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 (variation in CEO skill):

$$E_i[\tilde{y}_{it}] = E_i[\log y_{it} - \alpha \log w_{it}] = E_i\left[a_i + \varepsilon_{yit}\right] = a_i \implies Var(E_i[\tilde{y}_{it}]) = \sigma_0^2. \tag{2.24}$$

Parameters that drive the Board's decision The optimal wage depends on the Board's current beliefs about CEO skill (expected productivity), the degree of board capture, the probability of SOP failure give the wage, and an adjustment cost if the Board lowers the wage. These forces are reflected in the following approximation to the optimal log wage, which

illustrates clearly how estimable parameters drive variation in CEO pay:

$$\log w_{it} \approx \underbrace{\log(\alpha E_{bt}[A_t])}_{\text{Expected productivity}} - \underbrace{\log(1-\lambda)}_{\text{Board capture}} - \underbrace{\log\left(1 + \frac{\chi_B}{(1-\lambda)} \times pdf^{\text{U}}_{sit}(s_{it} \times w_{it}) s_{it}\right)}_{\Pr(\text{SOP fail})} + \underbrace{g(w_{it}, w_{it-1})}_{\text{Adjustment cost}},$$

where $pdf^{\text{U}}_{sit}(\cdot)$ is the likelihood function of the Shareholder's signal distribution.[27]

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

$$\log w_{it} = b_0 + b_1 \mathbf{1}[\text{SOP fail}_{it}] + b_2 \log w_{it-1} + \epsilon^b_{it}. \tag{2.25}$$

$\hat{b}_0$ maps to the average observed wage and is influenced by cross-sectional productivity and the Board's bias. The former is determined by identified output/skill parameters. Absent bias, (counterfactual) wages would solely reflect expectations of CEO skill: $\hat{b}_0$ thus identifies λ. $\hat{b}_1$ 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 (higher expected productivity) 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, $\hat{b}_2$ corners the dynamic aspect of the Board's problem (persistence in wage across periods) and identifies c_w.

Parameters that drive the Shareholder's decision The model's definition of SOP failure can be written closed-form as

[27] 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 $AC'(w_t, w_{t-1}) = -2c_w(w_t - w_{t-1})\mathbf{1}[w_t < w_{t-1}]$. Then, $g(w_t, w_{t-1}) \approx \log\left(2 + \frac{AC'(w_t, w_{t-1})}{1+\chi_B \times (1-\lambda)^{-1} s_{it} f(s_{it} w_t)}\right)$.

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

where $w_t^{\mathrm{U}}\left(p z_{sit} + (1-p) z_{yit}\right)$ is a linear combination of the two signals z_{sit} and z_{yit} transformed to the relevant distribution as in (2.13), and s_{it} is the Shareholder's choice variable.

The model-implied determination of SOP failure maps to the regression in the data

$$\mathbf{1}[\text{SOP fail}_{it}] = s_0 + s_1 \log w_{it} + s_2 \epsilon_{it}^{y} + \epsilon_{it}^{s}, \tag{2.26}$$

where ϵ_{it}^{y} is the residual from the output regression (2.23). $\hat{s}_0$ maps to the unconditional failure rate and is most informative for the parameters χ_S and λ. All else equal, higher χ_S will lower the observed rate of SOP failure and higher λ 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, which helps corner how much of SOP is driven by the private signal z_{sit}.

Bunching in Shareholder votes The observed bunching in SOP vote outcomes is informative about χ_B and χ_S. Before continuing, I define the data- and model-distances from the SOP failure threshold:

$$\Delta^{\text{data}} = \text{share against} - k, \quad \Delta^{\text{data}} \geq 0 \iff \mathbf{1}[\text{SOP fail}] \tag{2.27}$$

$$\Delta^{\text{model}} = F_{st}^{\mathrm{U}}(s_{it} \times w_{it}) - F_{st}^{\mathrm{U}}\left(Z_{syit}\right), \quad \Delta^{\text{model}} \geq 0 \iff \mathbf{1}[\text{SOP fail}], \tag{2.28}$$

where k is the empirical failure threshold (i.e., 30% or 50%). Focusing on the region close to the failure threshold (where bunching occurs) illustrates the identification argument:

matching the degree of bunching observed in the distribution of Δ^{data} in Δ^{model} helps to identify χ_B and χ_S.

In the model, the Shareholder commits to a single s_{it}, thereby making the realized threshold a function of the Board's signal z_{bit}. The Board has an incentive to *bunch* wage choices if the reduced expected cost of SOP failure outweighs the utility gained by a higher wage: bunching must be a function of the Board's cost to SOP failure χ_B. At the same time, χ_S dictates the location of the thresholds in the distribution function of Z_{syit}. A high χ_S will push the threshold toward zero and bunching will not occur; a low χ_S will push the threshold towards the middle of the distribution function and bunching similarly cannot occur.

It is useful to compare the model's notion of bunching and the bunching that occurs in Figure 2.2. In the data, this type of bunching is partially driven by large pivotal blockholders that can swing the outcome of a vote from a close fail to a close pass (as in Pinnington, 2022). There is no notion of blockholders in the model (just a single representative Shareholder). However, the *economic* force is the same: the blockholder (in the data) and the Shareholder (in the model) must have a low posterior belief about CEO skill for the vote to fail. The estimation targets a simple measure of bunching (Antill, 2021), which is similar to the estimator used in Figure 2.2, but easier to implement within the estimation; Appendix A.5.2 describes in detail.

2.4.2 Estimated Model Fit

Table 2.6 Panel A displays the closeness of auxiliary-model outcomes estimated on the observed and simulated data. The final column displays the test statistic from a two-way t-test comparing each outcome. Overall, the fit is quite good and importantly, the estimation matches the reduced-form outcomes which identify key structural parameters λ, χ_B and χ_S.

Outcomes which identify output and skill parameters are closely matched (outcomes

1-4). For the Board, average log wages (outcome 5, $\hat{b}_0$) and the log difference in wages in SOP failure (outcome 6, $\hat{b}_1$) are both matched well. The closeness in these outcomes is reassuring as they are key to identifying λ and χ_B. Persistence in log wages and variance of the wage regression residual (outcomes 7 and 8) are both over-estimated; implying that the estimation requires too-large an adjustment cost c_w and Board-signal volatility σ_{z_b}.

The estimation matches the observed SOP failure rate very closely (outcome 9): 7% and 6.7% in the observed and simulated data, respectively. The sensitivity of SOP failure to the wage (outcome 10, $\hat{s}_1$) is slightly under-estimated in the model, but there is no statistically significant difference. These are first-order reduced-form outcomes to match and the closeness is again reassuring of the estimation's success. The sensitivity of SOP failure to the output residual is well over-estimated in the model (outcome 11, $\hat{s}_2$). By incorporating the effect of company performance (through it being a signal of CEO ability) into the SOP vote outcome (Fact 1 and Fisch et al., 2018), it becomes too dependent on the productivity shock; this is natural given the model's parsimony. The simulated data matches the variance of the SOP regression residual (outcome 12) closely.

Lastly, bunching in shareholder votes (outcome 13) is matched closely in magnitude (0.121 *vs.* 0.116 in the estimation), though there is a statistically significant difference between the data and model. Bunching is very informative of the cost parameters χ_B and χ_S, so the closeness suggests these parameters are identified.

2.5 Results

2.5.1 Estimated Structural Parameters and Economic Implications

Table 2.6 Panel B displays the estimated parameters and quantifies their economic magnitudes.

Board Capture

The estimated value for λ is 0.612. I can translate this into how the Board splits the surplus between the Board and shareholders. In Appendix A.5.7, I show that the split of the surplus the Board wants to pay the CEO can be written as

$$\theta_{\text{CEO}} = \frac{\lambda w_t}{E_t[\pi_t] + \lambda w_t} = \frac{\lambda w_t}{E_t[A_t] w_t^{\alpha} - w_t + \lambda w_t}. \tag{2.29}$$

θ_{CEO} would equal zero without board capture and would equal one if the CEO captured all profits. Focusing on a CEO of average skill, I estimate $\theta_{\text{CEO}} = 40.7\%$. This is similar to Taylor (2013), who estimates that CEOs capture around 50% of surplus on the upside, but bear no downside risk. These estimates suggest that board capture is substantial at public companies. In Section 2.5.4, I show that estimated board capture varies with an empirical measure of CEO influence on the Board (board co-option, Coles et al., 2014).

The Board and Shareholder Costs to SOP Failure

I estimate that the Board considers SOP failure to be equivalent to 2.06% of firm value. For shareholders, this cost is 0.76% of value (or of a shareholder's equity stake), and both estimates are statistically different from zero.[28] It is important to note that these are utility costs. SOP failures do not affect output or 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.[29] Further, both costs are considerably lower in expectation as the unconditional failure rate is 6.7% in the simulated data: the Board (Shareholder) expected cost is closer to 0.14% (0.05%) of value.

[28] The parameters capturing the failure cost to the Board χ_B and Shareholder χ_S are estimated to be 0.240 and 0.088, respectively. To interpret these magnitudes, I normalize them by average firm value (in Appendix A.5.7, I show that average firm value V_0 can be derived in closed form).

[29] The exact source of the shareholder cost from SOP failure cannot be determined from the estimated model (see Fact 4 for possible sources of this cost).

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.[30] The threat of costly SOP failure disciplines the Board even when failure is unlikely. As I show in Section 2.5.2 providing shareholders with a punishment technology improves firm value considerably, *even though* punishing the Board is conditionally very costly.

Economic Implications of Other Parameters

The mean of CEO skill is normalized to zero (see Section 2.4.1). The volatility of CEO skill is 0.542, implying that a CEO at the 75th percentile of ability is about twice as productive as the 25th percentile.[31] This implies that CEO skill matters (supporting Bertrand and Schoar, 2003).

The standard deviations of the Board's and shareholder's private signals are 1.996 and 0.697, respectively. While the volatility of the innovation in the Board's signal is larger than the Shareholder's, it is important to remember that 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. The output shock volatility σ_y is 1.043. This implies that a large portion of observed variation in output comes from randomness, rather than CEO skill. Figure A.6 displays how the variance of (Board and Shareholder) beliefs decline over the CEO's tenure. While uncertainty about CEO ability decreases relatively quickly, there is still quite a lot of uncertainty through the median length of the CEO's tenure ($\tau = 7$, as in Table 2.1).

[30] 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).

[31] By CEO skill being normally distributed, the 25th (75th) percentile of productivity is ≈ 0.67 standard deviations below (above) the mean. Hence, the relative productivity of the 75th to 25th percentile CEO is $\approx \exp(1.34 \times \sigma_0)$.

2.5.2 How Much Does SOP Impact CEO Pay 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 as a governance mechanism impacted compensation policy and firm value.

Table 2.7 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 wages would increase if SOP were removed. It shows that, on average, SOP brings wages down by 4.4%, with considerable heterogeneity in the impact: the decrease in wages is 8.4% at the 75th percentile, with no change at the 25th. 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. The table also shows that SOP increases firm value by 4.6% on average. This is strikingly similar to the estimates from Cuñat et al. (2016), who find that the adoption of SOP increased market value by 5%.

This analysis benchmarks the key outcomes of the estimation against two important papers studying how the adoption of SOP affected compensation policy and value. At the same time, it shows that shareholder voice via a regularly-occurring opportunity to dissent from the Board on compensation policy is valuable to shareholders. Even though SOP failure is conditionally costly to shareholders, they do not have to use the punishment technology often in order to impact compensation policy and value considerably.

2.5.3 Model Estimation Validation

Untargeted Features of the Data

Confidence in the model's implications will be buttressed if the simulated and observed data yield similar results when examining outcomes not targeted in the estimation.

The relation between CEO pay and SOP vote outcomes In Figure 2.5, I explore the relation between CEO pay levels and SOP vote outcomes. For the observed and simulated data, I estimate a regression of log CEO pay on the outcome of the SOP vote (expressed as distance from the failure threshold, eqs. 2.27 and 2.28). The figure shows a striking similarity in the slope relation between CEO pay and the vote outcome, though the relation is tighter in the model.

SOP failure across the distribution of operating income Figure 2.6 Panel A displays a comparison of observed SOP failure rates in the observed and simulated data conditional on different levels of operating income. The estimation only targets the unconditional SOP failure rate and how the productivity shock affects SOP results, leaving untargeted how SOP results vary by operating income. The SOP failure rates in the observed and simulated data line up closely in the first three quartiles of operating income, but the relation deteriorates in the top quartile: 6.1% and 2.7% in the observed and simulated data, respectively.[32]

Changes in CEO pay following SOP disapproval Figure 2.6 Panel B displays how the observed and simulated data compare in terms of *changes* in CEO compensation in response to different SOP disapproval rates, or the relation between the SOP disapproval rate (relative to the failure threshold) and the change in CEO compensation from t to $t + 1$. The panel shows that the slope of this relationship is very similar in the model and data. 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 their future compensation decisions, even though the model does not directly incorporate how SOP failure influences future compensation policy decisions.

[32] The reason for the higher failure rate in the data in the top quartile of operating income is that CEOs in the top quartile receive higher compensation, both current and long-term compensation; these are the CEOs at the most productive firms and thus demand higher pay. Shareholders react to the level of CEO pay (see Table 2.2), which can help explain why the simulated data is different.

Is the Shareholder Cost to SOP Failure Necessary?

A key innovation of this paper is presenting evidence of shareholder cost to SOP failure and estimating the precise magnitude of this cost. Despite the empirical and survey evidence presented in Fact 4, direct evidence of the magnitude of the cost comes entirely from the structural model. In Table A.4, I estimate a reduced version of the model without the shareholder cost to SOP failure (i.e., $\chi_S = 0$) to highlight the importance of a positive shareholder failure cost. The data soundly reject this model. For example, the SOP failure rate (outcome 9) is too low: 7% and 4.6% in the observed and simulated data, respectively.

With no shareholder cost, shareholders are willing to fail the SOP in any state. To match the low failure rate, in this alternative specification, the model must push the degree of board capture down towards zero. With $\lambda = 0.092$, this implies that the Board wants to shift only 10% of surplus to the CEO (one fourth of the value from the main estimation and against previous research, such as Taylor, 2013). The Board SOP failure cost shoots up to 7.15% of firm value (relative to 2.06% in the main estimation). Taken together, these findings further imply that a positive shareholder cost to SOP failure is a key feature of the data.

2.5.4 Subsample Analysis

In this section, I use subsample analysis to explore how structural parameters and their implications vary with empirical measures of board capture (board co-option, Coles et al., 2014) and blockholder concentration (Hartzell and Starks, 2003). To accomplish this, I estimate the model on each subsample and display the resulting parameters from each split in tandem in Table 2.8, along with t-statistics from tests of difference in parameters.[33]

[33] Table A.5 displays the parameters (with standard errors) and fit for each subsample.

Board co-option

A key result of my main estimation is that there needs to be a high degree of board capture for the model to fit the data. While the model takes no stance on the channel through which board capture arises, Table 2.3 shows that *board co-option* (Coles et al., 2014), which measures the percentage of the Board appointed during the tenure of the CEO, plausibly plays a role. If CEOs influence director selection to tilt Board decisions in their favor, then estimated board capture should vary with empirical board co-option. Importantly, this would suggest the existence of the agency problem (CEO influence on the Board) which SOP is designed to mitigate, and that estimated board capture is not merely CEO bargaining power.

I split the sample by the degree of board co-option. To net out the mechanical increasing relationship between co-option and CEO tenure, I regress the degree of board co-option in my sample on CEO tenure fixed effects and save the residuals. As the model requires an entire CEO spell (from first year of tenure to separation), I take the average of the residualized board co-option across a CEO's tenure as my splitting variable. I sort firms into low (high) co-option if they are below (above) median average residualized co-option.

Table 2.8 columns 1-3 display the results of this subsample estimation. The model estimates a much larger degree of board capture in the high board co-option sample (row 1, 0.357 *vs.* 0.632). This means that low co-option CEOs capture 30.5% of the surplus produced by the match, and high co-option CEOs capture 51.3%. This evidence suggests that CEOs who have influence over the Board can distort compensation design and extract higher rents from the firm.

Interestingly, the Board cost is smaller in the high co-option sample, whereas the shareholder cost is larger (rows 11 and 12). In Table 2.3 Panel B, I show via reduced-form analysis that board co-option lessens the impact of SOP votes on future compen-

sation policy. My estimates suggest this is driven by the smaller Board cost and larger shareholder cost to SOP failure. This possibly suggests that powerful CEOs reduce the effectiveness of SOP as a governance mechanism via their influence on the Board.

Institutional ownership concentration

Institutional owners, by their size and consequent influence on corporate policies, often take on the role of disciplining management (Kakhbod et al., 2023; Appel et al., 2016; Duan and Jiao, 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. My model and estimates can help shed light on this apparent tension.

If large blockholders discipline compensation policy, the model predicts that the Board's SOP failure cost should be higher with the presence of large blockholders. The model also implies that the Shareholder SOP failure cost should increase with blockholder concentration. When the shareholder base is dispersed, no single investor is focal enough to take the brunt of a failed vote; large blockholders are more likely to feel the brunt of SOP failure. As such, the shareholder cost to SOP failure should also increase with blockholder concentration.

I use two measures of blockholder concentration, as inspired by Hartzell and Starks (2003). The first is the percentage of the market capitalization held by the top five institutional investors (top 5 inst. ownership). The second is the HHI of institutional investor base. Columns 4-6 of Table 2.8 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; columns 7-9 display the same for the HHI of the institutional shareholder base.

Both measures show that the Board's failure cost is higher when there is more con-

centration of institutional investors. For example, columns 4-5, rows 11-12 show that the Board cost is 2.6% (5.6%) 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 0.4% to 2.3% of firm value in the low vs. high sample. Interestingly, for both sample splits in the table, the degree of Board capture is not statistically significantly different, suggesting that this split hones in on how the SOP failure costs vary in the data.

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. The largest blockholders (passive funds) are not ineffective monitors *per se*, rather the negative consequences for going against the Board are higher.

2.6 Counterfactual Say-on-Pay Mechanism: Granting a Focal Shareholder an Advisory Board Seat

My analysis so far has taken as given the way SOP is implemented in practice. SOP is an *ex post* approval vote of realized compensation and not an *ex ante* advisory vote on proposed compensation, shareholders generally do not partake or share their views with the compensation committee during the wage-setting process. Thus, asymmetric beliefs about the CEO arise naturally and influence the outcomes of SOP votes, beyond any conflict of interest between the Board and shareholders. The efficacy of SOP could plausibly be improved if the Board and shareholders could coordinate their beliefs *in advance* of the compensation-setting process.

In this section, I alter the framework of my structural model to allow information-sharing between the Board and shareholders. This information-sharing emulates granting a focal shareholder a non-voting, advisory seat on the Board, a governance mechanism often put forth to align the beliefs (and hence actions) of management and share-

holders (Kakhbod et al., 2023).[34] This can also be interpreted as giving shareholders an *ex ante* say on *proposed* compensation policy, as opposed to an *ex post* opportunity to ratify the Board's decision.

In the following exercise, it is important to note that I keep the same set of structural parameters presented in the main analysis in Section 2.5: I am able to implement this counterfactual solely by changing the way information is revealed in the model. Figure 2.7 explains the intuition of the counterfactual. As shown in the top timeline, the baseline model inherently implies *belief disagreement*, as the Board and Shareholder signals are private information when strategies are set. In the counterfactual, B and S share their signals in advance of the vote, thus mitigating disagreement. Sharing of private signals means the SOP vote is solely a vote on performance *vs.* pay, as the vote is determined by the wage and productivity.

It should be mentioned that this counterfactual has its limitations. For example, investors state that they prefer not to vote against in SOPs because of the monitoring cost associated with SOP failure (repeated meetings with the compensation committee to discuss future compensation policy, Edmans et al., 2021). More stringent monitoring would imply a larger cost.

2.6.1 The Compensation Committee and SOP Vote

Like in Section 2.3.4, each period t the Board meets and receives its signal z_{bt} about CEO ability. However, now the Shareholder S shares its signal z_{st} with B and *vice versa*. Given beliefs $a \sim N(\mu_{at}, \sigma_{at}^2)$ at the start of t and denoting $\mathbf{z}_t = \begin{pmatrix} z_{bt} & z_{st} \end{pmatrix}$ as the shared signals,

[34] Granting a focal shareholder a non-voting board seat may resemble venture capitalists will often take on a "board observer" role, where they attend meetings and share their views, but hold no voting rights. See Kakhbod et al. (2023), Section 3.3.2 for a discussion.

the shared beliefs at the compensation committee will be

$$\mu_{at|\mathbf{z}_t} = \sigma^2_{at|\mathbf{z}_t}\left[\frac{\mu_{at}}{\sigma^2_{at}} + \frac{z_{bt}}{\sigma^2_{z_b}} + \frac{z_{st}}{\sigma^2_{z_s}}\right]$$

$$\sigma^2_{at|\mathbf{z}_t} = \left[\sigma^{-2}_{at} + \sigma^{-2}_{z_b} + \sigma^{-2}_{z_s}\right]^{-1}. \tag{2.30}$$

Now, the Board's wage decision is a function of both z_{bt} and z_{st} (i.e., $\mathbf{z}_t$). The Board's problem is very similar to (2.19), yet their beliefs are as in (2.30).[35] Because B and S share their signals, S no longer needs to integrate out B's private information, so they choose s_t to maximize

$$\max_{s_t} \quad E_{at|\mathbf{z}_t}[A_t]w_t(\mathbf{z}_t, s_t)^\alpha - w_t(\mathbf{z}_t, s_t) - \chi_S F^{\mathrm{U}}_{st}(s_t \times w_t(\mathbf{z}_t, s_t)) \tag{2.31}$$

The information sharing removes the possibility of belief disagreement between B and S: the operator $E_{a_t|\mathbf{z}_t}[\cdot]$ implies that B and S have the same beliefs when they set their strategies.

Two positive effects occur. First, there is a *cooperation effect*. By sharing their beliefs, B and S "meet in the middle," removing the possibility of disagreement over the CEO's skill: B and S arrive at a midpoint in cases when they would have very different beliefs, and in cases when their signals ultimately agree, the wage is set at a more appropriate level. Second, there is a *precision effect*, as beliefs mechanically become more precise: the Board's compensation policy converges faster to a wage appropriate for the CEO's true type.

I solve this version of the model using the parameters in Table 2.6. The counterfactual admits the same outcomes as the baseline, so I can compare observed quantities.

[35] Formally, the Board's Bellman equation is

$$V(\mu_{at}, \tau, w_{t-1}) = \max_{w_t} E_{at|\mathbf{z}_t}[A_t]w_t^\alpha - (1-\lambda)w_t - \chi_B F^{\mathrm{U}}_{st}(s_t \times w_t) - AC(w_t, w_{t-1}; \tau_t) + \delta_B\Big[(1-f(\tau_t))E_{at|\mathbf{z}_t}[V(\mu_{at+1}, \tau+1, w_t)] + f(\tau_t)V^R\Big]$$

The only difference to (2.19) is that the expectation operator follows from (2.30), as opposed to (2.18).

Table 2.9 displays the results: the counterfactual SOP failure rate and the average percentage changes in the CEO wage and firm value. The table also displays the distribution of percentage changes.

The SOP failure rate falls from 6.7% in the main estimation to 5% in the counterfactual. While the change is moderate, shareholder welfare increases as the punishment technology is used less often. Wages also fall on average, but this is not a given. When the Board and Shareholder agree that the CEO is good, wages rise (at the 75th percentile, the wage increase is 5%). On average, they fall because of the conflict of interest in the model. The Board's bias leads to suboptimal (from the shareholders' perspective) overproduction (empire-building). When the Board learns that the Shareholder thinks the CEO is low-type, this has a relatively stronger effect on decreasing the wage than when the Shareholder thinks the CEO is high-type, which occurs via the curvature of the production function. As such, the 25th percentile change in wages (10.8% decrease) is larger in magnitude than the 75th (5% increase).

Firm value increases on average by 4.9%, which is additive to the estimated impact of SOP from the baseline model (Section 2.5.2). At the lower end, firm value still increases considerably (at the 25th percentile, the increase is 3.7%). This is because wages fall on average, in most cases operating income is brought closer to its profit-maximizing level. Moreover, because of the precision effect, Board beliefs converge faster to the CEO's true ability, so a bad (good) CEO is less likely to be overpaid (underpaid) early on in tenure.

All in all, my model predicts that information-sharing between the Board and shareholders, actualized by granting a focal shareholder on the compensation committee, would positively impact firm value. However, the structure of the counterfactual implies perfect information-sharing by the Board and Shareholder. Partial information-sharing is more likely in practice, so the impacts on wages and firm value can be considered upper bounds. Further, advising on proposed compensation, via a non-voting Board seat or otherwise, may imply large monitoring costs.

2.7 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, provides shareholders with a platform to discipline the Board and curb CEO influence. As compensation arises as the key tool to mitigate the agency problem afflicting managerial decision-making, SOP votes are a potentially important governance mechanism. Yet, given their non-binding nature and low failure rate, the impact of SOP on compensation and value is unclear.

This paper establishes the impact of SOP votes and clarifies the economic mechanism by which this impact 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. These costs are equivalent to 2.06% and 0.76% of firm value for the Board and shareholders, respectively. The shareholder's cost of giving punishment explains the low failure rate of SOP, but does not mitigate its value-creation. I find that providing shareholders with a punishment technology through SOP reduces wages by 4.4% and increases firm value by 4.6% on average.

I also document what these costs are. SOP failure is a career and reputation concern for compensation committee directors. Shareholders internalize a cost to dissenting because this may damage the relationship with the Board and increase the likelihood of CEO turnover. Bunching in SOP votes provides objective evidence of these costs.

I use my estimates to show that SOP as a governance mechanism can be improved. I construct a counterfactual which emulates giving a focal shareholder a non-voting position on the Board, which induces sharing of private information. This counterfactual leads to a lower SOP failure rate, decreases wages and increases firm value on average.

My results provide insight about shareholder voting more generally, suggesting that high support rates in shareholder votes may not imply ineffective monitoring: the un-

observable (off-equilibrium) threat of a failed vote may have a large impact on corporate policies, even if this impact is not directly observable in the data. A policy implication is, if corporate decision-makers internalize a cost to failing environmental or social shareholder proposals, then a mandated, regular vote on these issues similar in spirit to SOP may positively impact shareholder and social welfare, even if these votes were to have a high observed pass rate.

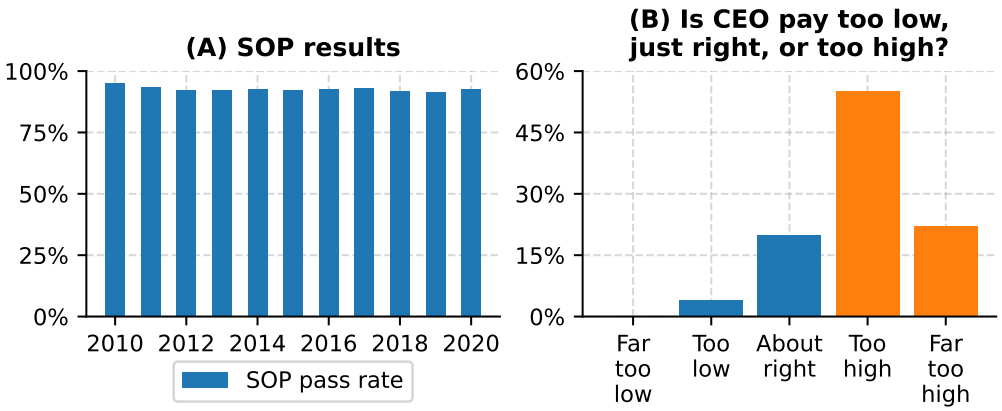

FIGURE 2.1: Say-On-Pay Results and Shareholder Satisfaction with CEO Pay

This figure displays motivation for the paper. Panel A displays SOP vote results in the US by year from 2010 to 2020; it shows that the percentage of SOP votes that pass (garner over 70% support, see Appendix A.2) is about 93%. Panel B displays survey data from Edmans et al. (2021), based on a question that asks UK institutional investors about the levels of the CEOs' pay; over 75% of survey respondents believe that their CEO is overpaid.

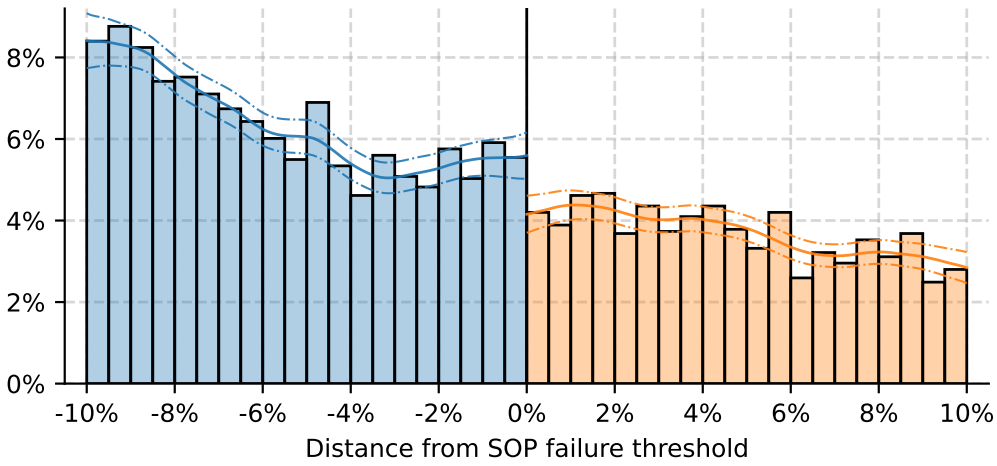

FIGURE 2.2: Density Manipulation of SOP Vote Outcomes

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 look for bunching at k in $\Delta^{\text{data}} = \text{share against} - k$, $\Delta^{\text{data}} \geq 0 \iff$ SOP fail. I focus on SOP votes falling within 10 pp of each failure threshold and test for density manipulation at zero. The blue and orange bars display observed frequencies of Δ^{data} in 0.5% bins and the blue and orange lines (and shaded areas, 95% confidence intervals) display the estimated density.

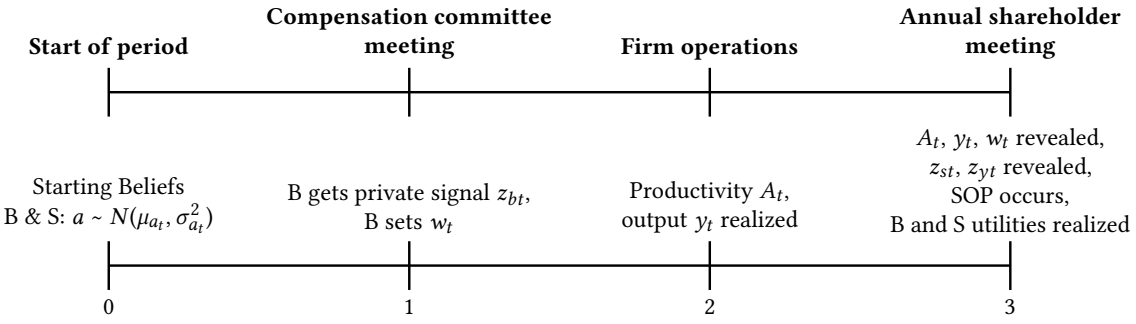

FIGURE 2.3: Model Timeline

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 A.5 displays an in-depth timeline which incorporates the timing of the strategies by the Board and Shareholder (See Appendix A.4.4).

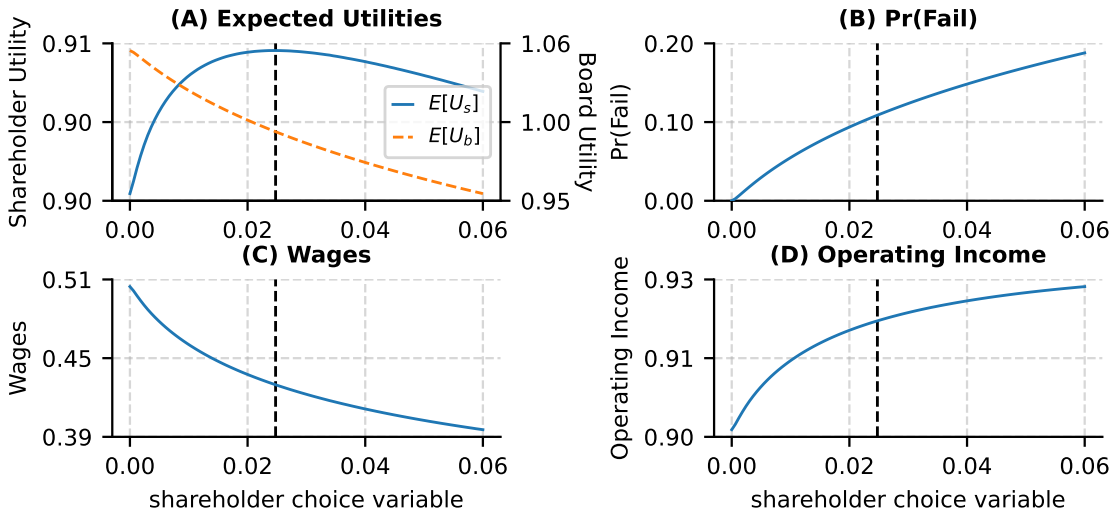

FIGURE 2.4: Comparative Statics Over Shareholder Strategy

This figure shows how shareholder strategy s influences per-period expected outcomes. To produce the figure, I fix $s \geq 0$ and produce the Board's best response for that s. Each plot shows expected values taken over the Board's signal z_b. For example, Panel A (blue line) displays shareholder utility averaged over the Board's signal z_b, $\int_{z_b} f(z_b)\big(E_s[A]w(z_b)^{\alpha} - w(z_b) - CDF_s^{\mathrm{U}}(s \times w(z_b))\big)dz_b$. The vertical, black and dashed line displays the equilibrium: S' best response.
Parameters: $\mu_0 = 0$, $\sigma_0 = 1$, $\sigma_y = 1.5$, $\sigma_{z_b} = 2.5$, $\sigma_{z_s} = 3$, $\chi_B = 0.45$, $\chi_S = 0.1$, $\alpha = 0.25$, $\lambda = 0.3$, $c_w = 0$

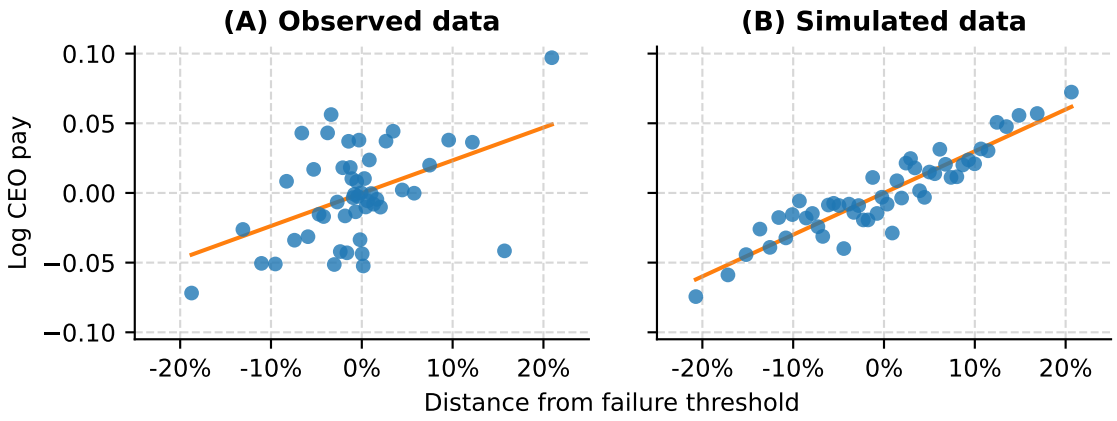

FIGURE 2.5: The Relation Between CEO Pay and SOP Vote Outcomes

This figure displays the relation between the level of CEO pay and the SOP vote outcome. In both panels, I display a binned scatterplot estimated from a regression of log CEO pay on SOP disapproval, expressed as the vote's distance from the failure threshold. Each regression includes CEO and CEO tenure fixed effects.

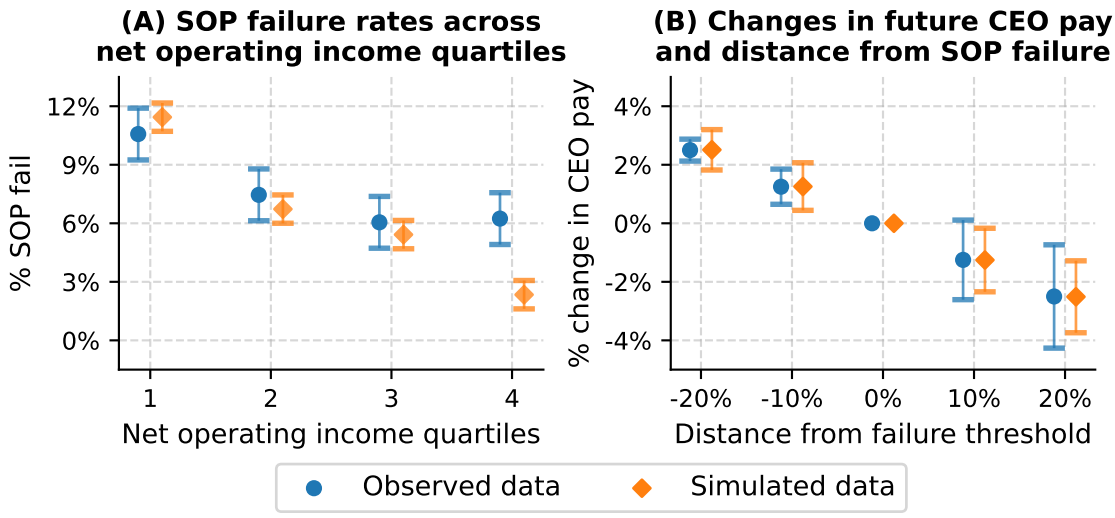

FIGURE 2.6: SOP Failure by Firm Performance and Changes in CEO Pay After SOP disapproval

This figure displays two model and estimation validation exercises. I simulate a dataset from the model using the parameters presented in Table 2.6. Panel A displays observed SOP failure rates across different quartiles of net operating income (industry-adjusted revenue less CEO compensation). Panel B displays estimated changes in CEO compensation in response to how far the previous year's SOP was from the SOP failure threshold. That is, I estimate the regression % Change in Compensation$_{i,c,t,t+1}$ = $\alpha + \beta$Distance from SOP failure$_{i,c,t}$ for firm i, CEO c and year t, and display the predicted change in compensation for different SOP vote outcomes. For both the model and data, I normalize the change in compensation to be 0% at vote failure.

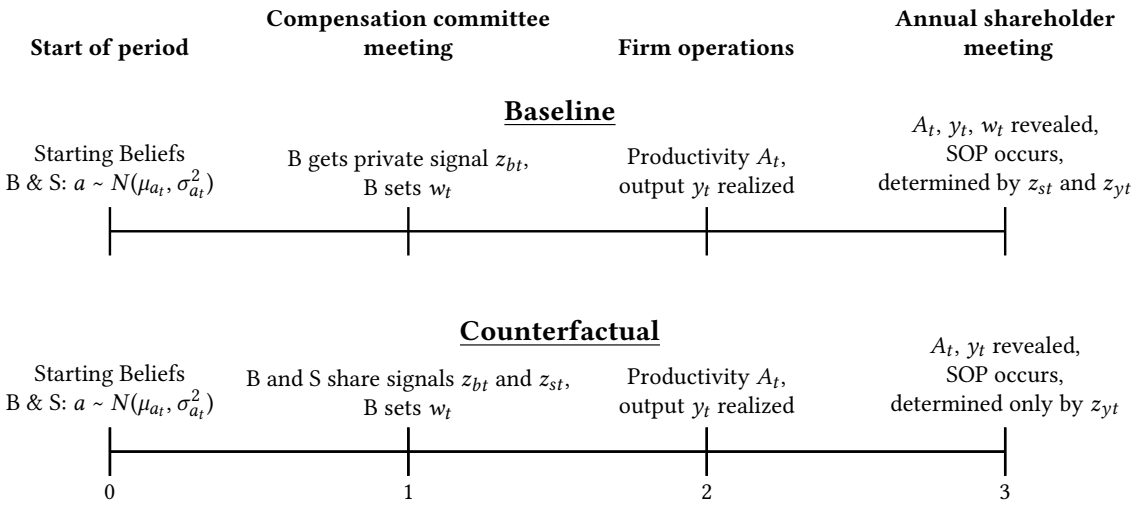

FIGURE 2.7: Sharing of Information: Counterfactual Model Timeline

This figure displays the within-period model timeline of the counterfactual, where shareholders have a non-voting position on the compensation committee. The top timeline ("Baseline") displays the timeline from the main version of the model (re-printed from Figure 2.3), the bottom timeline ("Counterfactual") displays the counterfactual timeline.

Table 2.1: Summary Statistics

	N	Mean	Std Dev	25%	50%	75%
Firm						
Assets ($b)	10,001	22.717	128.168	0.879	2.858	9.416
Revenues ($b)	10,001	7.703	21.147	0.596	1.749	5.612
Return on assets (%)	10,001	0.124	0.094	0.070	0.119	0.172
12-month stock return (%)	10,001	0.132	0.323	-0.056	0.113	0.293
Say-on-Pay						
% voting against in SOP	10,001	0.090	0.121	0.023	0.044	0.092
1[Less than 30% support]: **SOP failure**	10,001	0.068				
1[Less than 20% support]	10,001	0.129				
1[Less than 50% support]	10,001	0.018				
CEO						
Salary ($m)	10,001	0.855	0.383	0.600	0.808	1
Bonus ($m)	10,001	0.138	0.640	0	0	0
CEO tenure (years)	10,001	6.502	5.144	2	5	10
Length of CEO tenure (years)	2,528	8.311	5.449	4	7	12

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. I present statistics for the firm, SOP outcomes and the CEO. There are a total of 2,528 CEO spells in the dataset. I limit to CEO spells of greater than one year and at most 25 years.

Table 2.2: SOP Results, CEO Pay and Company Performance

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	SOP fail $\{0,1\}$				% vote no in SOP			
Log CEO compensation	0.015***	0.014***	0.019***	0.032***	0.007***	0.007***	0.010***	0.015***
	(0.004)	(0.003)	(0.006)	(0.011)	(0.002)	(0.002)	(0.003)	(0.005)
Stock return		-0.013***	-0.011***	-0.004		-0.010***	-0.009***	-0.007***
		(0.003)	(0.003)	(0.004)		(0.001)	(0.001)	(0.002)
Return on assets		-0.021***	-0.022***	-0.020**		-0.016***	-0.017***	-0.014***
		(0.006)	(0.007)	(0.008)		(0.003)	(0.003)	(0.004)
Log firm assets		0.039***	0.041***	0.030*		0.018***	0.018***	0.009
		(0.013)	(0.014)	(0.018)		(0.006)	(0.006)	(0.008)
Lagged log CEO compensation				-0.014*				-0.005
				(0.008)				(0.004)
Observations	9,841	9,841	9,556	6,736	9,841	9,841	9,556	6,736
R-squared	0.313	0.322	0.378	0.3c86	0.392	0.410	0.468	0.487
Firm FE	Yes	Yes			Yes	Yes		
Year FE		Yes	Yes	Yes		Yes	Yes	Yes
CEO tenure FE		Yes	Yes	Yes		Yes	Yes	Yes
Firm × CEO FE			Yes	Yes			Yes	Yes

This table explores the relation between SOP outcomes and CEO compensation, in support of Fact 1. Panel A estimates the relation between SOP outcomes and the level of CEO pay. The dependent variable in columns 1-4 is an indicator for SOP failure (at least 30% of shareholders voting against the SOP), and in columns 5-8 it is the percentage of shareholders that vote against. Log CEO compensation is the natural logarithm of the CEO's total current pay (salary and bonus). Panel B estimates 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. All covariates are defined in the appendix. Stock return 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%.

Table 2.3: Board Capture, CEO Compensation and SOP

Panel A: Board Capture and the Level of CEO Pay

	(1)	(2)	(3)	(4)
	Log CEO compensation			
Board co-option	0.081***	0.096***	0.095***	0.070***
	(0.029)	(0.024)	(0.024)	(0.023)
Return on assets		0.213***	0.213***	0.135***
		(0.023)	(0.023)	(0.024)
Stock return		0.020*	0.020*	0.036***
		(0.012)	(0.012)	(0.013)
Log firm assets		0.568***	0.577***	0.644***
		(0.023)	(0.030)	(0.033)
Log board size			-0.013	0.018
			(0.024)	(0.023)
Observations	8,865	8,865	8,865	8,865
R-squared	0.023	0.277	0.277	0.333
CEO tenure FE	Yes	Yes	Yes	Yes
Year FE				Yes
Industry FE				Yes

Panel B: Board Capture Modulates the Effect of SOP Disapproval on Changes in CEO Pay

	(1)	(2)	(3)	(4)	(5)	(6)
	Log change in CEO compensation					
Board co-option	0.014	0.010	0.006	0.015	0.016	0.012
	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)
SOP fail $\{0, 1\}$	-0.245***	-0.257***	-0.248***			
	(0.035)	(0.037)	(0.038)			
Board co-option × SOP fail $\{0, 1\}$		0.067*	0.063*			
		(0.037)	(0.037)			
% vote no in SOP				-0.078***	-0.081***	-0.079***
				(0.009)	(0.009)	(0.010)
Board co-option × % vote no in SOP					0.019**	0.019**
					(0.009)	(0.009)
Return on assets			-0.037***			-0.041***
			(0.008)			(0.008)
Stock return			0.056***			0.052***
			(0.009)			(0.009)
Log firm assets			-0.013			-0.007
			(0.013)			(0.013)
Log board size			0.002			-0.000
			(0.007)			(0.007)
Observations	6,388	6,388	6,388	6,388	6,388	6,388
R-squared	0.020	0.021	0.034	0.026	0.027	0.039
CEO tenure FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE			Yes			Yes
Industry FE			Yes			Yes

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. Panel B presents similar analysis to Table 2.2, with an interaction between board co-option and our two measures of SOP disapproval: an indicator for SOP failure (30% or greater vote against in the SOP), and the percent of votes against in the SOP. All covariates 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%.

Table 2.4: Evidence of Costs to Directors from SOP Disapproval

Panel A: Director Turnover

	(1)	(2)	(3)	(4)	(5)	(6)
	Director turnover {0, 1}					
SOP fail {0, 1}	2.303***	1.755***	1.512***			
	(0.537)	(0.504)	(0.513)			
% vote no in SOP				0.823***	0.673***	0.584***
				(0.176)	(0.168)	(0.173)
Stock return		-0.404**	-0.381**		-0.385**	-0.363**
		(0.174)	(0.173)		(0.174)	(0.174)
Return on assets		-1.460***	-1.625***		-1.443***	-1.603***
		(0.179)	(0.181)		(0.179)	(0.182)
Log firm assets		-0.354**	-1.032***		-0.386**	-1.033***
		(0.179)	(0.202)		(0.179)	(0.201)
Log board size			0.332**			0.342**
			(0.160)			(0.160)
Log director tenure			1.363***			1.364***
			(0.159)			(0.159)
Log director age			2.394***			2.391***
			(0.184)			(0.184)
Log CEO tenure			-1.199***			-1.204***
			(0.163)			(0.163)
Log CEO pay			0.567***			0.521***
			(0.150)			(0.151)
Constant	10.128***			10.418***		
	(0.173)			(0.162)		
Observations	33,213	33,213	33,213	33,213	33,213	33,213
R-squared	0.001	0.174	0.186	0.001	0.174	0.186
Year FE		Yes	Yes		Yes	Yes
Industry FE		Yes	Yes		Yes	Yes

Panel B: Compensation Committee Turnover, Conditional on Remaining on Board

	(1)	(2)	(3)	(4)	(5)	(6)
	Compensation committee turnover $\{0,1\}$					
SOP fail $\{0,1\}$	1.462***	1.126**	1.227***			
	(0.449)	(0.452)	(0.459)			
% vote no in SOP				0.640***	0.519***	0.572***
				(0.156)	(0.157)	(0.161)
Stock return		-0.481***	-0.470***		-0.460***	-0.445***
		(0.141)	(0.141)		(0.141)	(0.141)
Return on assets		-0.514***	-0.545***		-0.495***	-0.515***
		(0.148)	(0.152)		(0.148)	(0.152)
Log firm assets		0.097	-0.112		0.071	-0.107
		(0.153)	(0.192)		(0.153)	(0.192)
Log board size			0.326**			0.336**
			(0.140)			(0.140)
Log director tenure			0.193			0.195
			(0.145)			(0.145)
Log director age			-0.234			-0.240
			(0.152)			(0.152)
Log CEO tenure			-0.795***			-0.805***
			(0.140)			(0.140)
Log CEO pay			0.091			0.032
			(0.156)			(0.158)
Constant	5.460***			5.646***		
	(0.144)			(0.137)		
Observations	29,752	29,752	29,752	29,752	29,752	29,752
R-squared	0.000	0.008	0.009	0.001	0.008	0.010
Year FE		Yes	Yes		Yes	Yes
Industry FE		Yes	Yes		Yes	Yes

Panel C: Reduction in Outside Board Positions, Conditional on Remaining on Board

	(1)	(2)	(3)	(4)	(5)	(6)
	Reduction in outside board positions {0, 1}					
SOP fail {0, 1}	1.933**	1.879**	1.722**			
	(0.854)	(0.861)	(0.860)			
% vote no in SOP				0.488*	0.463*	0.398
				(0.262)	(0.265)	(0.266)
Log director tenure			-0.187			-0.192
			(0.274)			(0.274)
Log director age			1.935***			1.937***
			(0.290)			(0.290)
Log board size			0.203			0.200
			(0.257)			(0.257)
Outside firm average ROA			-0.126			-0.130
			(0.264)			(0.265)
Constant	9.486***			9.704***		
	(0.293)			(0.282)		
Observations	13,469	13,469	13,469	13,469	13,469	13,469
R-squared	0.000	0.002	0.006	0.000	0.002	0.006
Year FE		Yes	Yes		Yes	Yes
Industry FE		Yes	Yes		Yes	Yes

This table displays correlations between costly outcomes for directors and SOP disapproval rates. I focus on directors that serve on the compensation committee during the year of an SOP vote. In Panel A, director turnover 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. In Panel B, Compensation committee turnover occurs when a director is no longer on the compensation committee the year after an SOP vote, conditional on the director remaining as a member of the Board. In Panel C, a reduction in outside Board positions occurs when the number of outside Boards the 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 (i.e., it is an indicator variable). 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. All covariates are defined in the appendix; 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%.

Table 2.5: CEO Turnover and SOP Disapproval

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	CEO turnover $t+1$ $\{0,1\}$							
SOP fail $\{0,1\}$	3.169***	3.161***	2.458**	2.327**				
	(1.201)	(1.212)	(1.176)	(1.184)				
% vote no in SOP					1.142***	1.081***	0.806**	0.749*
					(0.377)	(0.382)	(0.383)	(0.386)
Log CEO compensation		-2.550***	-3.279***	29.524		-2.597***	-3.315***	29.119
		(0.629)	(0.684)	(24.560)		(0.631)	(0.689)	(24.553)
Stock return		-1.728***	-1.482***	-1.215		-1.695***	-1.459***	-1.200
		(0.349)	(0.353)	(1.183)		(0.349)	(0.353)	(1.181)
Return on assets		-1.197**	-1.654***	-0.277		-1.150**	-1.618***	-0.242
		(0.505)	(0.557)	(2.622)		(0.506)	(0.558)	(2.620)
Log firm assets		1.034***	1.136	1.168		1.059***	1.131	1.170
		(0.311)	(1.226)	(1.299)		(0.312)	(1.227)	(1.300)
Constant	9.435***				9.752***			
	(0.114)				(0.006)			
Observations	12,378	12,378	12,378	12,378	12,378	12,378	12,378	12,378
R-squared	0.123	0.131	0.214	0.216	0.123	0.131	0.214	0.216
Firm FE			Yes	Yes			Yes	Yes
CEO tenure FE			Yes	Yes			Yes	Yes
Year FE			Yes	Yes			Yes	Yes
Performance Polynomials				Yes				Yes

This table analyzes how CEO turnover changes with SOP disapproval. In all specifications, the dependent variable is a CEO turnover indicator in the year following a SOP vote. In columns 1-4, the main independent variable is an indicator for SOP failure, in columns 5-8 it is the percentage of shareholders who vote against the CEO's compensation in the SOP. All controls variables are from the year of the SOP. In columns 5 and 8, I include 5th polynomials of the firm's stock return and return on assets to control for possibly non-linear effects of performance on CEO turnover. Standard errors are displayed below coefficients and clustered at the firm× CEO level. ***, **, * denote significance at 1%, 5%, 10%.

Table 2.6: Structural Estimation Results

Panel A: Data and model moments

	Description	Notation	Data		
			Observed	Simulated	t-stat
(1)	Average log output	y_0	7.540	7.540	0.001
(2)	CEO-average output variance	$Var(E_i[\tilde{y}])$	0.880	0.829	1.000
(3)	Elasticity of output to wage	y_1	0.832	0.797	0.620
(4)	Output residual variance	$Var(\epsilon^y)$	1.278	1.260	0.235
(5)	Average log wage when SOP passes	b_0	-0.140	-0.137	-0.972
(6)	Change in log wage when SOP fails	b_1	0.042	0.055	-0.838
(7)	Log wage persistence	b_2	0.794	0.873	-6.466
(8)	Wage residual variance	$Var\left(\epsilon^b\right)$	0.059	0.077	-5.683
(9)	SOP failure rate	s_0	0.070	0.067	0.674
(10)	SOP failure—wage sensitivity	s_1	0.061	0.038	1.826
(11)	SOP failure—output residual sensitivity	s_2	-0.019	-0.034	2.456
(12)	SOP failure residual variance	$Var(\epsilon^s)$	0.064	0.063	0.164
(13)	Bunching estimator	$\boldsymbol{B}$	0.121	0.116	3.909
		χ^2 (p-val)	65.296 (0.000)		

Panel B: Parameter Estimates

Description	Notation	Value
Parameters		
CEO board capture	λ	0.612 (0.006)
Board SOP failure cost	χ_B	0.240 (0.007)
Shareholder SOP failure cost	χ_S	0.088 (0.024)
Prior average of CEO ability	μ_0	0
Prior std dev of CEO ability	σ_0	0.542 (0.006)
Output—CEO wage elasticity	α	0.263 (0.011)
Std dev of productivity shock	σ_y	1.043 (0.010)
Std dev of Board signal	σ_{z_b}	1.996 (0.025)
Std dev of Shareholder signal	σ_{z_s}	0.697 (0.014)
Scaling factor	$\log \eta$	7.572 (0.011)
CEO wage adjustment cost	c_w	5.201 (0.373)
Estimated quantities		
CEO surplus capture	θ_{CEO}	40.70% (0.84%)
Board SOP cost (% average value)	χ_B / V_0	2.06% (0.12%)
Shareholder SOP cost (% average value)	χ_S / V_0	0.76% (0.23%)

This table contains outcomes of the estimation procedure. Panel A displays model and data moments, as well as the difference between each moment. In Panel A, I perform a t-test of the difference in each model and data moment; t-statistics are displayed in the last column. In the last row of the table, I display the result of the χ^2 test of overidentifying restrictions. Panel B displays estimates of the parameters that drive the model in Section 2.3. As explained in Section 2.4, μ_0 is normalized to 0, hence this parameter is not estimated. Parameters are estimated using SMM, which is described in detail in Appendix A.5. The panel also displays the degree of CEO surplus capture implied by the board capture, and the magnitudes of the Board and Shareholder costs to SOP failure as a percentage of the model's 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 A.5.7 for a derivation).

Table 2.7: The impact of SOP on Compensation Policy

	Mean	**25%**	**50%**	**75%**
SOP failure rate	—			
Percent change in				
Wages	+4.35%	+0.00%	+2.72%	+8.40%
Firm value	-4.58%	-6.22%	-4.47%	-3.40%

This table analyzes the impact of SOP on the Board's pay decisions. I simulate a counterfactual 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. As SOP no longer enters the model, there is no counterfactual SOP failure rate to display. Rows 2 and 3 display the counterfactual percentage change in wages and firm value when removing the Board cost.

Table 2.8: Subsample Heterogeneity

		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
		Board co-option			**Top 5 inst. ownership**			**HHI inst. ownership**		
	Notation	**Low**	**High**	***t*-stat**	**Low**	**High**	***t*-stat**	**Low**	**High**	***t*-stat**
Parameters										
(1)	λ	0.36	0.63	-25.79	0.54	0.51	0.95	0.59	0.58	0.08
(2)	χ_B	0.41	0.28	5.12	0.40	0.46	-4.16	0.45	0.79	-7.17
(3)	χ_S	0.04	0.18	-1.48	0.07	0.19	-2.43	0.09	0.36	-2.34
(4)	σ_0	0.50	0.36	7.99	0.61	0.43	3.69	0.58	0.62	-2.06
(5)	α	0.43	0.32	8.10	0.28	0.37	-3.80	0.31	0.33	-2.43
(6)	σ_y	1.04	0.92	8.68	1.19	0.94	8.34	1.08	0.93	4.53
(7)	σ_{z_b}	0.73	0.98	-10.90	1.06	0.76	3.17	2.85	0.32	28.49
(8)	σ_{z_s}	1.08	2.22	-7.42	0.64	2.25	-11.39	0.51	3.70	-4.36
(9)	$\log \eta$	7.75	7.76	-0.41	8.08	7.17	21.47	8.31	6.94	32.92
(10)	c_w	4.02	3.23	3.68	5.67	4.12	4.64	5.91	4.65	2.12
Estimated quantities										
(11)	θ_{CEO}	30.5%	51.3%	-9.99	33.6%	41.5%	-4.40	44.0%	45.6%	-0.14
(12)	χ_B / V_0	4.2%	3.4%	3.83	2.6%	5.6%	-7.00	3.7%	8.0%	-13.06
(13)	χ_S / V_0	0.4%	2.2%	-1.62	0.4%	2.3%	-2.95	0.7%	3.7%	-3.03

This table contains outcomes of the estimation on the main sample split by different characteristics. Each subsample split is estimated using the same routine as the main sample (see Table 2.6). Following Section 2.5.4, “Board co-option" measures the percentage of directors appointed during the CEO’s tenure, with CEO tenure fixed effects residualized out (as co-option increases mechanically in tenure). Following Section 2.5.4, “Top 5 inst. ownership" calculates the percentage of firm’s equity held by the five largest institutional investors and the HHI of the institutional ownership base, both are measures of institutional ownership concentration (Hartzell and Starks, 2003). I take the average of these measures across a CEO’s tenure and split the sample into below/above median, so low (high) refers to below (above) the median taken cross-sectionally across CEO spells. I display the parameters from the estimation on each split, along with the t-statistics from a test of equality of the parameters from the split. For example, columns 1 and 2 show the estimated parameters from splitting the sample into above- and below-median Board co-option (Coles et al., 2014), and column 3 displays the t-statistic from testing parameter equality. Columns 1-3 focus on board co-option. Columns 4-6 focus on percentage of equity held by the top five largest institutional investors, and columns 7-9 focus on concentration in the institutional shareholder base (as measured by the HHI).

Table 2.9: Counterfactual Exercise: Granting a Focal Shareholder an Advisory Board Position

	Mean	**25%**	**50%**	**75%**
SOP failure rate	4.97%			
Percent change in				
Wages	-2.05%	-10.84%	-2.51%	+5.01%
Firm value	+4.87%	+3.68%	+4.19%	+5.07%

This table displays a counterfactual experiment where I re-solve a different version of the model, in which the Board and shareholders are allowed to share their signals before the compensation contract is decided (i.e., the shareholder is given a non-voting, advisory seat on the Board). The first row displays the counterfactual SOP failure rate. Rows 2 and 3 display the counterfactual percentage change in wages and firm value. To compute these changes, 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. I then display the average percentage change, along with the 25th, 50th and 75th percentiles.

3

Project Development with Delegated Bargaining: The Role of Elevated Hurdle Rates

"If I set a higher hurdle rate, my people go out and get better deals on contracts they negotiate."

CFO of NYSE listed energy firm

3.1 Introduction

Project *development* is distinct from project *evaluation*. By the latter we mean the capital budgeting steps prescribed by textbooks, while the former is broader and involves interactions with trading partners, including the procurement of inputs. In this paper, we explore the special role of discount rates in project development, and in particular, the common practice of requiring that expected returns for projects exceed a hurdle rate that is greater than the risk-appropriate weighted-average cost of capital (WACC). We show that, because of the process through which projects are cultivated, adding a buffer to discount rates often preserves and improves firm value.

Two ideas are central. First, every project that a firm accepts is an acquisition that changes the boundary of the firm. Each one needs to be developed. Therefore, up-front

investment costs are not exogenous, contrary to what most textbooks posit. They are the *endogenous* result of a collection of bargaining outcomes between a firm's managers and outsiders who do business with the firm. If sufficient gains through trade exist, investment in positive NPV projects creates a joint surplus, over which the firm and its trading partners negotiate.

For example, suppose that building a new plant requires land (which is part of up front investment costs C_0). The price that the land sells for will depend on what it will eventually be used for (the value created by future cash flows) and the relative bargaining power between the firm's managers and the owners of the property.

Second, employees accept their CFO's instructions about hurdle rates as given without questioning or verifying them.[1] This common organizational structure within firms gives rise to a delegated bargaining situation, in which top executives can alter the firm's observable preferences to outsiders.[2] As such, when a hurdle rate is high, project managers presume that their walkaway value is higher and that the firm has better outside options. This can result in the firm obtaining a greater share of the surplus during negotiations as projects are developed.

Based on these ideas, we pose a theoretical model of delegated bargaining, characterize the optimal hurdle rate buffer in several settings, and test the novel empirical implications that arise from the analysis.

Hurdle rate buffers are extremely common. Our survey evidence shows that 78% of companies set hurdle rates above their cost of capital, and among these firms the average buffer is 6.6%.[3] The surveys also indicate why (Figure 3.1): because of perceived or real financial constraints (Graham, 2022), idiosyncratic risk (Décaire, 2021), uncertainty about

[1] Evidence based on CFO interviews from Graham (2022) illustrates that the hurdle rate is considered "sacred" at firms, thus providing a clear benchmark to facilitate decisions by mid-level employees.

[2] See Schelling (1956), Crawford and Varian (1979), Sobel (1981), Jones (1989), Fershtman et al. (1991), Burtraw (1993), and Segendorff (1998).

[3] Jagannathan et al. (2016), Gormsen and Huber (2022), and Graham (2022) provide similar estimates.

the true cost of capital (Bessembinder and Décaire, 2021), agency problems within the firm (Harris et al., 1982; Harris and Raviv, 1996, 1998; Chen and Jiang, 2004), managerial time constraints (Jagannathan et al., 2016), or for real options effects (McDonald, 2000).[4]

But, in many cases textbook capital budgeting with inflated hurdle rates would lead companies to destroy value ("leave money on the table") as they bypass positive net present value (NPV) projects.[5] In these cases, why wouldn't innovation in contracting, governance, hiring, or organizational structure arise to claw back lost firm value? And why wouldn't the market for corporate control, discipline from earnings reporting, or stock market pressure correct this behavior? A key outcome of our model is showing that inflated hurdle rates persist because they are part of an equilibrium trade-off between the gains to trade resulting from negotiating with an inflated hurdle rate and losses from bypassing moderately positive NPV projects. The result of this trade-off can on net preserve firm value and at times create value.

Importantly, the purpose of this paper is not to propose a "7th explanation" for the use of hurdle rate buffers or to run a horse race among all of the explanations. Rather, we seek to explore the effects that elevated hurdle rates have on project development and why buffers appear to persist over time, whether buffers arise for bargaining power (strategic) reasons or because of traditional (non-strategic) reasons.[6] As such, our rationale for the persistence of buffers co-exists with and complements the reasons already described in the literature.[7] Nonetheless, as we show, our model and empirical evidence

[4] For reasons to be explained below, we refer to these collectively as "non-strategic" explanations of the hurdle rate buffer.

[5] The magnitudes in our surveys imply that firms impose a negative bias of 5.7% for one-year cash flows and a 45% decrease for perpetual cash flows. Among firms that use a hurdle rate buffer, the average hurdle rate and WACC are 14.73% and 8.15%, respectively. Each dollar of perpetual cash flow is worth $6.79 in present value under the hurdle rate, and $12.27 under the cost of capital. This represents a decrease in value of 45% attributable to the buffer.

[6] If buffers were value-destroying, we would expect them to have disappeared.

[7] We repeat this point for emphasis: Even if firms do not "strategically" choose a high hurdle rate to aid negotiations, but instead the hurdle rate is high for traditional "non-strategic" reasons, the key point of our paper still goes through: a high hurdle rate aids bargaining. Thus, our message does not compete with nor

provide implications that are not consistent with what other explanations for buffers would predict.

In our model of delegated bargaining, we derive the Nash bargaining solution from first principles. A key result is that parties enjoy a greater share of the surplus when their required rate of return is reported to be higher. This is in the spirit of Rubinstein (1982), but distinct. Rubinstein bargaining uses a higher discount rate to model less patience, so a party would be vulnerable to waiting to make a deal: a higher discount rate is associated with a lower split of the surplus (Rubinstein and Wolinsky, 1985). Here, the opposite occurs because the value of the outside option is based on the opportunity cost of capital. When a party is required to earn a higher return, their perceived walkaway value increases and they earn more surplus if a deal is on the table.

So, in the model, while projects with moderate NPV may be bypassed when using an inflated hurdle rate, overall this can still create value due to the bargaining advantage the firm enjoys when negotiating the projects that the firm does undertake. This arises both when the CFO uses a buffer strategically or non-strategically (e.g., based on the explanations in Figure 3.1), and when a trading partner uses a buffer as well. This helps explain why hurdle rate buffers are so widely used and persistent.

Using data from six CFO surveys, we establish new stylized facts about hurdle rates and test the model's predictions about buffer use. One interesting new empirical fact is that the distribution of ex post return on investment capital (ROIC, which we estimate using Compustat) aligns with the ex ante hurdle rates that CFOs provide in our surveys. Interestingly, a portion of the mass of the ROIC distribution is shifted from just below the hurdle rate to just above the hurdle rate, consistent with firms working to "meet or beat" the hurdle rate benchmark. These facts corroborate the validity of the hurdle rates that CFOs provide to us via surveys; and more broadly, they indicate that buffered-up

reject traditional rationales; rather, our argument complements traditional hurdle rate buffer rationales by offering an explanation for how large, long-lasting buffers might continue to exist.

hurdle rates are a focal point of corporate action. We additionally show that realized ROIC is typically several percentage points greater than WACC, indicating that hurdle rate buffers play a deeper role than simply offsetting optimistic cash flow projections provided by managers. These findings motivate our theoretical modeling of the effects of elevated hurdle rates

We use the survey data to empirically explore four specific predictions from the model on both the extensive and intensive margins. Our first two empirical tests relate the buffer to the firm's cost of capital and how much bargaining power a firm has. First, we confirm empirically that buffer use and the size of the chosen buffer are decreasing in a firm's cost of capital. While this is implied by our model, it differs from what one might expect based on buffer explanations in the prior literature. For example, if managers use a larger buffer to compensate for idiosyncratic risk (Décaire, 2021), we would expect there to be, on average, a positive association between the cost of capital and the buffer, because idiosyncratic risk is correlated with systematic risk, empirically. Also, prior explanations based on financial or managerial constraints would also generally predict either no relation between the buffer and cost of capital (because managerial time constraints are likely independent of the cost of capital) or a positive relation (because financial constraints are more likely to bind for high cost of capital firms). Thus, our prediction and result are unique in the literature.

Second, we show that buffer use is negatively correlated with plausible measures of a firm's bargaining power relative to trading partners. This is consistent with the model's implication that firms that already have strong bargaining power have less incentive to use buffered-up hurdle rates to further enhance bargaining power. We confirm the expected relation using a measure of the concentration of a firm's customer base (Patatoukas, 2012); confirming this negative relation is new to the literature. In a separate test of this prediction, among firms operating in predominantly B2B industries, for which markups based on Baqaee and Farhi (2020b) would seem to be a reasonable proxy for

bargaining power, we confirm the negative relation.

A third implication comes from the fundamental premise of the model: project value is a function of bargaining over inputs and outputs. In the data, we show that companies with high asset tangibility or that operate in industries where projects require more property, plant, and equipment (PPE) have higher buffer use. This positive relation is not obvious from explanations such as financial constraints or idiosyncratic volatility, as these quantities are typically negatively correlated with tangibility (Almeida and Campello, 2007; Fink et al., 2010). Under the view that, all else equal, bargaining will play a more important role in acquisitions of significant fixed assets compared to other inputs, the PPE findings are consistent with predictions from our model. But we note that there are other possible explanations for the positive relation we document. For example, to the extent that asset tangibility is related to lumpy investment dynamics (due to adjustment costs or some other friction, e.g., Cooper and Haltiwanger, 2006), we would expect to see the same directional relation. We validate our hypothesis relating to PPE by finding a negative relation between asset redeployability (Kim and Kung, 2017) and the buffer – all else equal, assets with lower redeployability (assets that are more specific-to-use) are more likely to be bargained over, hence increasing the incentive to use a buffer in negotiations.

A final implication of our model is that the buffer is larger for firms with higher volatility of underlying project value (i.e., idiosyncratic risk), consistent with Décaire (2021). We also show that buffer use increases in uncertainty about WACC (Bessembinder and Décaire, 2021). These findings tie our paper to frontier research on corporate discount rates.

The rest of the paper is organized as follows. In the rest of this section, we describe the contributions of our paper in the context of the existing literature. In Section 3.2, we describe the survey and other data that we use, and establish a number of stylized facts about the hurdle rate buffer. In Section 3.3, we analyze delegated bargaining theoretically

and characterize the empirical implications about IRR buffers that arise from the model. The proofs are contained in the appendix. In Section 3.4, we develop a set of hypotheses and test the model using data from six CFO surveys. Section 3.5 concludes.

3.1.1 Literature Review

The discrepancy between hurdle rates and WACC has been appreciated for decades. Starting with Poterba and Summers (1995), the authors surveyed CEO's from Fortune 1000 companies and showed that hurdle rates often exceed both the equityholders' average rates of return and the cost of debt. Many of the firms in their study were in manufacturing, where long-term capital budgeting is commonplace. Unlike our study, though, they did not ask their subjects whether a different WACC was also computed and not used for project development. While the authors did not compute a specific IRR buffer as we do, their findings suggest that the hurdle rates used were indeed inflated.

Since that time, many explanations have been proposed for the use of elevated hurdle rates. Harris and Raviv (1996) provide a theory of internal capital allocation in which there is decentralized information about projects and agency problems within the firm. The goal of headquarters is to manage the tendency for a division to over-invest in new projects. This may arise because divisional managers can employ less effort if more capital is invested (Harris et al., 1982), or because managers are either optimistic or have private benefits when projects are undertaken. Harris and Raviv (1996) posit that a mixture of capital constraints and oversight can help ameliorate these frictions.

Bernardo et al. (2001) investigate these issues in the context of contracting and show that only high-quality projects get funded, and that managers of these receive greater incentive compensation. This stems from information asymmetries between CFOs (headquarters) and divisional managers. Chen and Jiang (2004) show that asymmetric information is not necessary to cause use of higher hurdle rates. They show that use of an IRR buffer solves an agency problem in which the divisional manager is required to exert

costly, non-contractible effort to collect information. In both cases, headquarters must commit to the allocation and compensation schemes. Chen and Jiang (2004) surmise this could arise from the rigidity of the capital budgeting process. The analysis that we provide in our theoretical model supports this claim – and furthermore shows that such commitment that arises when the CFO (headquarters) dictates the elevated hurdle rate and project managers take it as given without verification can lead to bargaining benefits and higher overall valuations.

More recently, other explanations have been offered for the pervasive use of IRR buffers. Jagannathan et al. (2016) posit that IRR buffers arise because of real (or perceived) managerial constraints. They demonstrate this to be the case using survey data. They also quantify the wedge between hurdle rates and the firm's WACC. Their magnitudes are consistent with the results from our surveys, which is reassuring. Further, our more recent surveys confirm that elevated hurdle rates remain important for managers. However, as we describe above, one still wonders why firms would leave value on the table if they could access labor markets and hire or train more managers. Our paper addresses this issue directly, both theoretically and empirically.

Décaire (2021) shows that firms use an IRR buffer when they face higher idiosyncratic risk. Bessembinder and Décaire (2021) show that uncertainty about discount rates (systematic risk) causes an upward bias in estimated NPVs, leading to higher corporate investment; such a bias should cause firms to adjust their hurdle rates. These predictions are also considered in our model. We also show that they are present in the data, confirming these studies as well as predictions from our model; we control for these forces when investigating the other implications of our model.

Relatedly, a recent literature has studied the real effects of managers using imprecise or over-simplified discount rates. Krüger et al. (2015) analyze the practice of using a single firm-wide discount rate (see Graham and Harvey, 2001a) and show that companies value high-risk projects using discount rates that are too low and vice versa; this leads

to over-investment (under-investment) in relatively risky (safe) projects. They further show that this practice leads to lost value – in the context of acquisitions, when the bidder's beta is lower than the target's beta, announcement returns are significantly lower. Dessaint et al. (2021) study a different mechanism – managers and the market may arrive at different acquisition valuations if managers (over-)rely on CAPM-estimated discount rates. This stems from the well documented observation that the empirical security market line is too "flat" relative to CAPM-implied estimates. They show that managers who rely on the CAPM tend to overvalue low-beta targets and undervalue high-beta targets relative to the market. Together, these papers highlight additional real consequences of the use of imprecise WACC in corporate finance, which we incorporate into our model and empirical analysis.

Our paper is also related to Gormsen and Huber (2022), who document time-varying wedges between discount rates and the cost of capital disclosed in earnings calls. Despite our having much different data sources, the magnitude of the hurdle rate buffers are similar between our paper and theirs. Gormsen and Huber (2022) also show that the hurdle rate wedges are negatively related to investment, providing new evidence on the real effects of inflated discount rates.

A key contribution of our paper relative to the prior work on the buffer is that we show that a buffer can actually be value-enhancing. McDonald (2000) posits a very different mechanism that relates buffer use to increases in value. In that paper, an IRR buffer serves as a heuristic that approximates solving a real options problem. Managers have the option to wait to start projects, and thus choose when to optimally exercise an American option. If the option is sufficiently in the money, they do not wait; this is akin to the project return exceeding a high hurdle rate. Even if buffers are used heuristically in this manner, our model suggests that large buffers have the added benefit of impacting the bargaining surplus the project will achieve.

As discussed briefly above, our paper adds to the extant bargaining literature. The

model explored in Rubinstein (1982) and Rubinstein and Wolinsky (1985) is sequential in which two parties take turns making take-it-or-leave-it offers. Bargaining takes place over time, as long as an agreement has not been reached. Each party has a discount rate, which proxies for their eagerness to get to a negotiated deal faster. The solution involves an equilibrium split whereby a higher discount rate yields lower surplus. Our model yields the opposite intuition: having a higher discount rate increases a party's (perceived) walkaway value because the outside option is more valuable.

The organizational structure in our model gives rise to decentralization in which the CFO reports the hurdle rate and the manager develops projects. This is in the spirit of Schelling (1956), Crawford and Varian (1979), and Sobel (1981) who explore how advantage can be gained by distorting the impression or beliefs of a counterparty during negotiation. Explicitly, this can be done via delegation in which a party commits to using a representative without the ability to renegotiate suboptimal outcomes (Jones, 1989; Fershtman et al., 1991; Burtraw, 1993; Segendorff, 1998). The structure of our theoretical construct shares this feature, which we use to model typical organizational behavior.

Finally, our paper adds to the rent-seeking and innovation literature. Nash bargaining in our model yields a surplus share that resembles a classic Tullock contest function (Tullock, 1980), where the fraction of the project surplus the firm receives is an increasing function of the firm's hurdle rate. While this arises endogenously, it is natural as it has been used often in the rent-seeking literature to represent the proportion of the market each party enjoys when new projects or markets arise (Hirshleifer, 1989). Tullock contest functions are commonly used to characterize R&D races (D'Aspremont and Jacquemin, 1988; Chung, 1996; Andrei and Carlin, 2022). What differentiates our work is that the Tullock contest function arises from delegated bargaining and is a function of the purported cost of capital for each firm. Additionally, the purported hurdle rate also creates spillovers for each firm, in that positive NPV projects may be discarded. We thus add to prior work in the rent-seeking literature that considers spillovers where the size of the

pie in contests either increases with effort (D'Aspremont and Jacquemin, 1988; Chung, 1996) or shrinks (Alexeev and Leitzel, 1996).

3.2 Data Description and the Hurdle Rate Buffer

In this section, we describe our data sources, discuss the reasons that firms set their hurdle rates above their cost of capital, and describe variation in the hurdle rate buffer across observable dimensions. We discuss the main variables covered by our CFO surveys in Section 3.2.1, discuss the origins of the buffer in Section 3.2.2 and relate ex ante hurdle rates to ex post returns in Section 3.2.2. Table B.1 provides variable definitions and Appendix B.3 provides more detail on the CFO Survey.

3.2.1 CFO Survey Data

Our primary data source comes from six CFO surveys conducted by Duke University. These surveys have been conducted quarterly for decades; the most recent surveys have been conducted jointly with the Federal Reserve Banks of Richmond and Atlanta. At several points in time, the survey asked CFOs directly for both their hurdle rate and their weighted average cost of capital.

One advantage of gathering data via surveys is that we obtain information directly from each firm's primary financial decision-maker. Another advantage is that we are able to gather data via questions that precisely define hurdle rates and other variables versus having to infer or approximate the variables. We ascertain the firms' weighted average cost of capital by asking CFOs "what do you estimate is your firm's overall weighted average cost of capital (WACC)?" To obtain data on hurdle rates, we ask "what is your firm's hurdle rate (the rate of return that an investment must beat in order to be adopted)?" Similar questions appeared in each of the six CFO surveys that we analyze and the consistent wording of these questions gives us confidence that our measures capture what we in-

tend.[8] For sample inclusion, we require that the CFO supply a value for their WACC and hurdle rate. From there, we compute the hurdle rate buffer as the difference between the hurdle rate and WACC.[9]

Figure 3.2 summarizes key characteristics of the survey firms. Panel A shows that firms in our sample are distributed across several industries and that manufacturing firms comprise the largest portion of the sample. Panel B conveys that the sample includes large firms (revenue greater than $1 billion), as well as smaller firms. Panel C shows that our sample is comprised of both private and public firms. Panel D shows that our sample is relatively evenly distributed across six different surveys that asked CFOs about their hurdle rates.

Table 3.1 Panel A displays summary statistics of variables related to the hurdle rate and cost of capital. The average hurdle rate in our sample is 13.88% and the average WACC is 8.77%, giving rise to an average buffer of just over 5%. Nearly 78% of companies use a hurdle rate buffer. Among the firms that have a non-zero buffer, the average buffer is 6.6%. The size of the buffer is consistent with other research (Jagannathan et al., 2016; Graham, 2022; Gormsen and Huber, 2022).

3.2.2 Why Do Companies Use a Hurdle Rate Buffer?

In this section, we establish stylized facts about why firms use buffers and the cross-sectional and time series variation in the intensive and extensive margins of the buffer.

[8] The surveys are 2011q1, 2012q2, 2017q2, 2017q3, 2019q1 and 2022q2. Appendix B.3 provides detail on the data collection process and the questions posed to CFOs concerning their hurdle rate and cost of capital.

[9] In our terminology, we use the extensive margin to refer to the use of a positive buffer (i.e., an indicator taking one (zero) if the firm's hurdle rate is greater than (equal to) WACC). We use intensive margin to refer to the difference between the hurdle rate and WACC. Much of our analysis of the intensive margin is for the full sample (including firms that have buffer = 0), which is the same sample used for the extensive margin analysis and therefore enhances comparability; the intensive analysis is robust to using a Tobit specification (to account for zero-inflation in our intensive margin variable). Our intensive analysis is also robust to deleting buffer = 0 firms and examining only firms that have a positive buffer (e.g., see Panel B of Table 3.3). Table B.1 displays variable definitions.

Why Do CFOs Set Hurdle Rates Above the Cost of Capital?

Various frictions may lead a company to set its hurdle rate above the cost of capital. Companies may use a buffer to ration capital, so that they choose just the highest expected returns projects of a given class. This may arise due to constraints on funding, managerial time, non-managerial labor availability, or production capacity (Jagannathan et al., 2016; Graham, 2022). A second set of forces may lead CFOs to set a hurdle rate equal to their perceived cost of capital plus a "fudge factor," which serves to offset complications from idiosyncratic risk or difficulties in estimating the true cost of capital (Décaire, 2021; Bessembinder and Décaire, 2021), because of agency issues or adverse selection problems within the firm, or because of behavioral forces that lead managers to use overly-optimistic cash flow forecasts as they pitch their projects to upper management (Harris et al., 1982; Harris and Raviv, 1996, 1998; Chen and Jiang, 2004). Third, using a buffer may approximate decision-making that accounts for real options motivations (McDonald, 2000).

To explore these possibilities, on three of the CFO surveys we asked CFOs "Why do you set your hurdle rate above WACC?", allowing respondents to choose among a list of reasons representing the forces mentioned above (or to write in other explanations).[10] The available choices varied somewhat between surveys but always included choices related to five of the traditional explanations mentioned above: (i) financing constraints (ii) managerial/resource constraints (iii) project prioritization (iv) idiosyncratic risk/uncertainty and (v) over-optimism/agency. We classify CFO responses concerning why their hurdle rate exceeds WACC under these five qualitative groupings, in order to ascertain the relative importance of each in explaining the buffer.

Figure 3.1 presents the proportion of CFOs that selected reasons that fall within each of the five categorizations of why companies use hurdle rate buffers. Table B.2 describes

[10] The three surveys were conducted in 2011q1, 2019q1 and 2022q2.

how we group responses across each survey. Panel A of Figure 3.1 displays the importance of each category across the three surveys. There is reasonable support for each of the five categories. For example, 60% of respondents indicate that they implement a buffer due to project prioritization and more than half indicate that risk or uncertainty in estimating the cost of capital or unmeasured risks lead to buffers. Panel B displays the importance of each reason by survey year. The importance of financing and resource constraints dipped in 2019 when the economy was in a relatively strong position (relative to 2011, where the after-effects of the 2008 crisis were still present; and relative to 2022, with a tight labor market and rising borrowing costs). To emphasize a point made earlier, whether the CFO increases the hurdle rate buffer for one of the "non-strategic" reasons reflected in Figure 3.1, or to gain strategic negotiating advantage, the bargaining implications derived in our model still hold.

Panel C displays the results conditional on firm size (large firms being those with revenue weakly greater than $1 billion). Financing and resource constraints are relatively more important for small firms. Interestingly, idiosyncratic project risk or uncertainty about WACC or discount rate or other unmeasured risks are more important for large firms. This may be because large firms are more likely to have numerous diverse projects, with different idiosyncratic risk, and are more likely to use equity markets or issue debt, potentially adding complexity to estimating the cost of capital and the return required by investors.

Taken together, the information in Figure 3.1 helps us create empirical proxies to measure these buffer explanations, both to explore these explanations directly and to control for these reasons when we explore empirical predictions from the model. Section 3.4.1 discusses specific variables used to measure buffer explanations, as well as the other non-CFO survey variables that we use in our empirical analysis.

Variation in Hurdle Rate Buffers Across Firms, Time, and Geography

Surveys back to the 1980s show that firms' hurdle rates are on-average relatively constant across time – suggesting that buffers also retain this consistency, or are even growing through time (Graham, 2022; Sharpe and Suarez, 2021; Jagannathan et al., 2016; Poterba and Summers, 1995). To ascertain the degree of variation in the buffer in our dataset, Figure 3.5 investigates the cross-sectional and time series properties of the hurdle rate, the cost of capital, and the buffer. Panel A displays industry average hurdle rates, costs of capital and buffers. The black error-caps on each bar display the within-group inter-quartile range.[11] It is noteworthy that there is substantial variation *within* each industry, which is consistent with our model in which a firm's position within its industry affects its incentives to use hurdle rate buffers. Also clear is that the buffer is not just an artifact of fixed differences across industries, which will play an important role in our empirical analysis. Leveraging our data's fine industry classifications, Table B.3 displays average hurdle rate, cost of capital, and the intensive and extensive margins of the buffer across NAICS-2 and CFO survey industries.[12]

Panel B shows that the size of the average buffer is relatively stable across survey years. Panel C shows that there is a clear difference in the size of the buffer for small versus large firms. For firms with less than $1 billion in sales revenue, the average buffer is 5.8%, whereas for firms with revenue above $1 billion, it is 3.8%. Though not shown in the figure, among firms with a non-zero buffer, the average buffer for small/large firms is 6.9%/5.3%.

Panel D shows that the buffer is remarkably similar across regions of the world, despite variation in the hurdle rate and the cost of capital (see also Gormsen and Huber,

[11] For example, the average buffer in Mining/Construction is approximately 6%, and the inter-quartile range is 0% to roughly 9%.

[12] We use company names and the Infogroup dataset to match survey firms to their NAICS codes. Survey industry codes are supplied by the CFO directly and roughly equate to one-digit SIC codes.

2022). We separately examine the properties of the extensive margin of the buffer in Figure B.2. The extensive margin patterns in the appendix figure are very similar to the just-discussed intensive margin results shown in Figure 3.5.

Ex Ante Hurdle Rates versus Realized Returns

Companies use hurdle rates as a benchmark against which expected returns of projects are compared. To the extent that realized and expected returns align, we would therefore expect hurdle rates to influence the distribution of realized project returns. In particular, as we explain in the next two paragraphs, we expect to observe excess mass in the ROIC distribution directly above the benchmark hurdle rate. We do not observe project-level returns in our data; however, for public firms in Compustat we can calculate the return on invested capital (ROIC) as a measure of realized returns, which is an aggregation of the firm's return on all its projects.

The spirit of our model, and the data, suggest that hurdle rates are chosen to be greater than the true cost of capital. Thus, in project development, even with a bargaining advantage, managers of marginal projects (those with expected returns close to or just below the hurdle rate) might need to stretch to surpass the elevated hurdle. Managers are incentivized to push a little harder to get these marginal projects over the line. This could occur via the specific mechanism in our paper (bargaining hard over project inputs), though it could also result from another aspect of project development.

A direct empirical consequence of the just described evaluation-development link would be bunching of ROIC just above the hurdle rate; that is, a shift in a portion of the ROIC distribution from just below to just above the hurdle rate. Empirically, we search for this discontinuity in the distribution using a simple application of manipulation testing, often used in the context of regression discontinuity designs (McCrary, 2008; Cattaneo et al., 2018). In our application, we are interested in determining formally whether there is evidence of a discontinuity in the density of ROIC at the hurdle rate. To implement this

test, we define the variable $Excess\ ROIC = ROIC - Hurdle$, which measures the firm's ROIC earned in excess of its hurdle rate. We are interested in testing for the existence of a discontinuity in excess ROIC at zero.

Figure 3.4 displays the density of excess ROIC centered at zero. The light blue and orange bars display observed frequencies of excess ROIC in 0.5 percentage point bins. The frequency of observations just above zero (realized ROIC just greater than the hurdle rate) is much larger than that directly below. For example, about 9% of relevant observations are in the +0.5% bin, whereas about 3.5% are in the −0.5% bin.[13] The shaded area and lines overlaid on the histograms display the result from the methodology described in Cattaneo et al. (2018) – estimating the density of excess ROIC via local polynomial density estimation. The figures reveals a distinct discontinuity in excess ROIC at zero.[14]

The findings in Figure 3.4 provide evidence consistent with companies choosing projects to beat their hurdle rates. The bunching just above the hurdle rate is consistent with the spirit of our model. Moreover, if companies are choosing projects based on hurdle rates, rather than the cost of capital, one would not expect to find ROICs bunched just above the WACC. This is exactly what we find in Figure B.3, where we undertake a similar exercise to Figure 3.4, testing instead for discontinuity in the density of ROIC at firms' costs of capital. We find no statistical or visual evidence of bunching in ROIC directly above WACC, suggesting that firms target their hurdle rates, as opposed to WACC.[15]

[13] As stated in the caption of Figure 3.4, we linearly interpolate between observed hurdle rates for Compustat firms in our sample. For example, if the hurdle rate is 10% in 2012q2 and 11% in 2012q4, we assume a hurdle rate of 10.5% in 2012q3. Given the stickiness of hurdle rates as shown in Graham (2022), we feel that this assumption is reasonable. We confirm the primary findings of our analysis in Figure 3.4 by alternatively using "stair-step" interpolation, where we keep the hurdle rate constant between consecutive observation until data from a future survey reveals a change in the hurdle. For the same example just given, we would assume that the firm's hurdle rate remained at 10% in 2012q3. Results are very similar.

[14] We follow the data-driven bandwidth selection methodology to estimate the density, and set the order of the local polynomial to 3. Inference using this methodology can be sensitive to researcher-chosen parameters, in particular the order of the local polynomial and the size of the bandwidth for local estimation. In Table B.4, we vary these parameters and show that the implications of Figure 3.4 are quite robust to parameter choices, though the effects attenuate as the size of the bandwidth increases.

[15] In appendix Figure B.3, we use the same sample of firms as in Figure 3.4. In nearly all of our analyses, we

One take-away from these figures is explicit confirmation, for the first time in the literature to our knowledge, of the importance of hurdle rates (more so than the cost of capital) in project development, as well as "beat-the-benchmark" bunching just above the hurdle rate.[16] These implications motivate our focus on the hurdle rate in our model, as explored in the next section.

The evidence in this section confirms three empirical features of hurdle rate buffers. First, firms use buffers for a variety of reasons, often to ration project choice because of constraints, or to counteract frictions in the firm's organizational structure or capital budgeting process (Figure 3.1). Second, while average buffers are relatively constant across observable dimensions, there is considerable variation, confirming that the buffer is not an artifact of differences in financial practice across, e.g., industries (Figure 3.5). Third, the hurdle rate matters for realized returns, which is consistent with hurdle-based project evaluation influencing outcomes in project development (Figure 3.4). These three points motivate and discipline the model in Section 3.3 and subsequent empirical analyses in Section 3.4.

3.3 The Model

3.3.1 Preliminaries

Consider an unlevered firm that employs a CFO who is tasked with calculating a hurdle rate for the firm's activities. We assume that the CFO of firm F uses an asset pricing

require the respondent to supply both a hurdle rate and a WACC. However, for Figure 3.4 we maximize our sample size to include firms that only supply their hurdle rate. For this expanded sample, because CFOs do not always provide a WACC number along with their hurdle rate, in Figure B.3 only we estimate WACC using Compustat data and the CAPM (see Appendix B.3). We do this to facilitate comparison between the role of the hurdle rate vs. the role of WACC in determining the distribution of ROIC, as depicted in the two figures.

[16] Burgstahler and Dichev (1997) find beat-the-benchmark behavior for earnings per share targets, a highly visible external benchmark. Our finding shows similar behavior for the hurdle rate benchmark, a number which is internal to the firm and not generally known publicly, suggesting strong beat-the-benchmark corporate incentives.

model that computes the cost of capital $W_F = 1 + r^F_{wacc}$. The CFO reports a hurdle rate H_F to other employees in the firm, which is taken at face value and is not verified. As such, the CFO has discretion to report the real $H_F = W_F$, or she may use an IRR buffer and report a higher gross hurdle rate $H_F = \tau_F$ to others in the firm.

Reasons for reporting $\tau_F > W_F$ may be strategic or non-strategic. We refer to an IRR buffer as strategic if a positive buffer is chosen by the CFO to take advantage of the delegated-bargaining organizational structure of the firm (to be described shortly). Alternatively, IRR buffers may exist for non-strategic reasons that may be the result of compensating for a lack of precision in estimating W_F, conservatism, or managerial and financial constraints.

Employees within the firm who identify and manage projects are called *delegates*. The firm invests a fixed amount of capital x_F in each project. For any given project, we begin by considering that it is necessary for the firm to acquire one asset from an outside business partner (e.g., land, a building, or equipment). If the delegate fails to secure the necessary asset, the entire project becomes infeasible. In Section 3.3.2, we generalize this to consider N outside entities/assets.

The outside business partner, O, also has a CFO and a delegate who represent its interests. The business partner invests x_O in the asset and has a required return. In the same way, though, the delegate for the outside business partner is either endowed with their true cost of capital $H_O = W_O = 1 + r^O_{wacc}$ or a gross hurdle rate that includes an IRR buffer $H_O = \tau_O$.

The projects that the primary firm has access to are heterogeneous. Assume that the values from these investment opportunities are uniformly distributed across the support $V \sim U[0, \overline{V}]$. Given that x_F and x_O are fixed, there exists a unique R associated with each project. Assume that both delegates and CFO's have full information about this distribution. Also, to abstract away from other agency or adverse selection problems, assume that the value of any project that is ultimately accepted is ex post verifiable.

When assets are acquired, delegated Nash bargaining takes place, with a fraction θ of the surplus being allocated to the firm, F, and $(1-\theta)$ going to trading partner, O. Both delegates report their respective hurdle rates to each other, H_F and H_O. For each potential trade, Nash bargaining maximizes

$$\max_{s} \quad (s_F - d_F)^{\alpha}(s_O - d_O)^{\beta}, \tag{3.1}$$

where $s \equiv \{s_F, s_O\}$ is the surplus allocated to each party and $d \equiv \{d_F, d_O\}$ is each party's disagreement payoff.

Lemma 1. *For any feasible R, suppose that $s = \{\theta(x_F+x_O)R - x_F H_F, (1-\theta)(x_F+x_O)R - x_O H_O\}$, $d = \{0,0\}$, and $\alpha = \beta = 1$. Then,*

$$\theta = \frac{1}{2} + \frac{x_F H_F - x_O H_O}{2(x_F + x_O)R}. \tag{3.2}$$

In Lemma 1, the total value (V) equals the return on investment R times the total capital invested ($x_F + x_O$). Firms receive their portion of V minus their opportunity cost. If the two managers walk away, they believe that they can invest their capital in an alternative at their hurdle rate, earning zero NPV. The solution in (3.2) shows that the split of the surplus depends on each firm's purported hurdle rate - the firm with a higher hurdle rate effectively has more bargaining power. The two parties split the surplus, with a greater proportion going to the party with the higher walkaway value: θ is increasing in H_F and decreasing in H_O.

The use of discount rates in θ is in the spirit of Rubinstein bargaining, but distinct. Because Rubinstein bargaining is sequential, the party with less patience (higher discount rate) earns a lower split of the surplus. Here, the opposite occurs since the perceived value of the outside option is based on the hurdle rate. When a party is required to earn a higher return, their walkaway value increases, and they earn more surplus if a deal is feasible.

Having shown that bargaining will depend on the hurdle rates of each firm, we formulate a more convenient bargaining split that will be used in the rest of the paper, which

still yields an outcome with similar properties. We define $\alpha = x_F H_F$ and $\beta = x_O H_O$, and assume that the surplus to be split is the value of the project above and beyond what both firms would require if they made investments at their hurdle rates. Total net surplus is computed as $(x_F + x_O)R - (x_F H_F + x_O H_O)$. As before, walking away means that the firms earn a zero NPV.

Lemma 2. *For any feasible R, suppose that* $s = \{\theta[(x_F + x_O)R - (x_F H_F + x_O H_O)], (1-\theta)[(x_F + x_O)R - (x_F H_F + x_O H_O)]\}$, $d = \{0, 0\}$, $\alpha = x_F H_F$, *and* $\beta = x_O H_O$. *Then,*

$$\theta = \frac{x_F H_F}{x_F H_F + x_O H_O}. \tag{3.3}$$

According to Lemma 2, the firm and its trading partner can increase their respective bargaining power by reporting a higher hurdle rate. Equation (3.3) is increasing in H_F and decreasing in H_O, just as in Lemma 1. In the equilibrium that we derive below, this may make it attractive for both parties to use IRR buffers. As we will see however, the offsetting cost will be that positive NPV projects are lost because of the exaggeration of reported hurdle rates.

Interestingly, the expression in (3.3) resembles a 2-firm, logit Tullock contest function (Tullock, 1980), where each party gains a share of the value from the project (Hirshleifer, 1989). This is natural here as a bargaining payoff because Tullock contest functions are commonly used to characterize the gains from innovation in the rent seeking literature (D'Aspremont and Jacquemin, 1988; Chung, 1996; Baye and Hoppe, 2003; Andrei and Carlin, 2022).[17]

In the following subsections, we present four different versions of the model linking hurdle rates (and hurdle rate buffers) to project outcomes via this bargaining mechanism. We illustrate how this effect holds whether buffers are used by a single firm strategically,

[17] Chowdhury and Sheremeta (2011) generalize Tullock contest functions to consider linear combinations of effort complementarities in duopoly contests. See also D'Aspremont and Jacquemin (1988), Chung (1996), and Alexeev and Leitzel (1996).

by both firms competitively, or by a single firm for reasons unrelated to bargaining, either due to imprecise estimation of the WACC or due to operational or financial constraints. In all versions of the model, the hurdle rate buffer affects bargaining outcomes and firm value.

3.3.2 Strategic Use of IRR Buffers

Firms may use a buffer for a variety of reasons. In this section, we start by characterizing the buffer that maximizes firm value if it is used purposefully by the CFO to increase bargaining power. This will serve as a benchmark to compare cases where firms utilize a buffer for non-strategic reasons, such as constraints on managerial time or capital.

Single-firm IRR Buffer

It is instructive to start by assuming that the outside business partner does not use an IRR buffer strategically. Without loss of generality, we set $H_O = W_O$.[18] Nash bargaining takes place over each project value $V = (x_F + x_O)R$. From (3.3), if the firm does not use an IRR buffer, its split of the total surplus is given by

$$\theta_N = \frac{x_F W_F}{x_F W_F + x_O W_O}. \tag{3.4}$$

Firm value is computed as

$$V_N = \int_{\underline{V}}^{\overline{V}} \theta_N \, dV,$$

where $\underline{V} \equiv x_F W_F + x_O W_O$. As such, $\underline{V}$ measures the minimum value a project yields so that both parties can receive sufficient surplus to participate. Simplification yields

$$V_N = \frac{x_F W_F \overline{V}}{x_F W_F + x_O W_O} - x_F W_F. \tag{3.5}$$

[18] A value $H_O > W_O$ that is set exogenously or used for non-strategic reasons does not qualitatively change the results that follow.

Now consider that a hurdle rate buffer may be used strategically by the firm when the outsider does not. The firm solves

$$\max_{\tau_F \geq W_F} \quad V_B = \int_{\underline{V}'}^{\overline{V}} \theta_B dV, \tag{3.6}$$

where $H_F = \tau_F$ is the buffered (gross) hurdle rate and

$$\theta_B = \frac{x_F \tau_F}{x_F \tau_F + x_O W_O}.$$

We restrict $\tau_F \geq W_F$ because it is a dominated strategy to choose a lower cost of capital than W_F; in that case the firm would reduce its bargaining power and accept negative NPV projects. So, when a positive IRR buffer is used, $\underline{V}' \equiv x_F \tau_F + x_O W_O$, which is larger than $\underline{V}$. The lower limit of the integral increases because some positive NPV projects are rejected and the firm's delegate walks away.

So, the CFO faces a trade-off. Using a higher buffer increases the split of the surplus from θ_N to θ_B, but raises the minimum acceptable project from $\underline{V}$ to $\underline{V}'$: Bargaining power increases, but positive NPV projects are discarded.

Proposition 1. *For bargaining as in (3.3), when the firm uses a positive IRR buffer and the outsider does not, the optimal buffered hurdle rate is given by*

$$\tau_F^* = \frac{\sqrt{\overline{V} x_O W_O} - x_O W_O}{x_F}, \tag{3.7}$$

and the resulting value to the firm is given by

$$V_B^* = \overline{V} + x_O W_O - 2\sqrt{\overline{V} x_O W_O}. \tag{3.8}$$

The change in value resulting from use of the buffered hurdle rate is strictly positive, that is,

$$\Delta V = V_B^* - V_N = \frac{x_O W_O \overline{V}}{x_F W_F + x_O W_O} + (x_F W_F + x_O W_O) - 2\sqrt{\overline{V} x_O W_O} > 0. \tag{3.9}$$

ΔV is increasing in $\overline{V}$ and decreasing in x_F and W_F.

We can use (3.7) to determine a condition that characterizes when a positive buffer is used (extensive margin), that is when $\tau_F > W_F$. A positive buffer is used if

$$(1 - \theta_N)\overline{V} > x_F W_F + x_O W_O. \tag{3.10}$$

The firm is more likely to use the buffer when its starting bargaining power (i.e., their initial split of the surplus, θ_N) is lower, its own cost of capital (W_F) is lower, and when there are more high potential project values ($\overline{V}$). The upper bound of the project values $\overline{V}$ can also be considered a proxy for the uncertainty faced by managers, since the variance of the distribution increases with $\overline{V}$. If (3.10) is not satisfied, then the CFO reports the hurdle rate as equal to the true WACC.

As a specific example, Figure 3.5 displays the value implications from the model. We calibrate parameters so the optimal hurdle rate buffer $b^* = 5.11\%$, the unconditional average value in the CFO Survey data. On the right y-axis, we plot the change in value implied by a standard capital budgeting model, where we only consider the value lost from foregoing positive NPV projects. At b^*, the implied "traditional" change in value is about −4.4%. On the left y-axis, we incorporate the value benefit from bargaining and plot the model-implied percentage increase in firm value from using a hurdle rate buffer.[19] This equals +0.11%. While this may seem small, it does imply that firm value is preserved even with a large buffer. The difference in value creation between traditional assumptions and the bargaining model is quantitatively large at 4.5 percentage points.

Also according to Proposition 1, if an IRR buffer is used, the intensive margin implications are that buffers are more attractive for high potential project values and less attractive for higher firm investment and when there is a higher cost of capital. The relationship is more subtle between ΔV and $x_O W_O$. This stems from a non-monotonic relationship

[19] The difference between the two axes is the bargaining split. On the left y-axis, θ_B is allowed to change with b, so $V_B = \theta_B \overline{V} - (W_F + b)$ and the bargaining split increases with b. On the right, we hold the surplus split fixed, so $\tilde{V}_B = \theta_N \overline{V} - (W_F + b)$, with θ_N displayed in (3.4).

between the use of the IRR buffer (τ_F^*) and $x_O W_O$. Taking the following derivative

$$\frac{\partial \tau^*}{\partial W_O} = x_O \left(\frac{\overline{V}}{2x_F \sqrt{\overline{V} x_O W_O}} - \frac{1}{x_F} \right)$$

demonstrates that the relationship is positive only if $\overline{V} > 2\underline{V}'$. The same is true for the derivative with respect to x_O. When the surplus gained from a larger share of projects is greater than the value lost from forgoing positive NPV projects, the firm's IRR buffer is increasing in $x_O W_O$. Otherwise, if $\overline{V} < 2\underline{V}'$, the cost of lost positive NPV projects dominates and the buffer is smaller.

Competitive IRR Buffers

Now consider that each player $i \in \{F, O\}$ chooses an IRR buffer to maximize

$$\max_{\tau_F} \quad V_F = \int_{\underline{V}''}^{\overline{V}} \theta dV$$

and

$$\max_{\tau_O} \quad V_O = \int_{\underline{V}''}^{\overline{V}} (1 - \theta) dV,$$

conditional on their counterparty acting optimally as well. Now, the lower limit of integration is defined as $\underline{V}'' = x_F \tau_F + x_O \tau_O$, which is the purported minimum project value that can meet both parties' required returns under the buffered IRR hurdle rates.

Proposition 2. *For bargaining as in (3.3) there exists a unique symmetric Nash equilibrium in which the optimal buffered hurdle rates are given by*

$$\tau_F^* = \frac{\overline{V}}{4x_F} \qquad \tau_O^* = \frac{\overline{V}}{4x_O} \tag{3.11}$$

and the value for each firm is

$$V_F^* = V_O^* = \frac{\overline{V}}{4}. \tag{3.12}$$

Compared to the case in which hurdle rate buffers are infeasible, the deadweight loss is

$$Loss = \frac{\overline{V}}{2} - (x_F W_F + x_O W_O). \tag{3.13}$$

The comparative statics implied by Proposition 2 are straightforward. For each counterparty, the hurdle rate buffer is increasing in the potential value of projects ($\overline{V}$) and decreasing in the scale of each party's investment.

According to Proposition 2, hurdle rate buffers are used by both firms as long as $\overline{V} > 4x_F W_F$ and $\overline{V} > 4x_O W_O$. The condition in (3.10) provides a sufficient condition for this to hold for firm F. A similar condition can be calculated for firm O as

$$\theta_N \overline{V} > x_F W_F + x_O W_O. \tag{3.14}$$

While there is an equilibrium deadweight loss, this is not necessarily a Prisoner's dilemma. Let us compare each party's value when both use an IRR buffer to the case where neither (are allowed to) do so.

$$\Delta V_F = \overline{V}\left(\frac{1}{2} - \theta_N\right) - x_F(\tau_F - W_F) \tag{3.15}$$

$$\Delta V_O = \overline{V}\left(\theta_N - \frac{1}{2}\right) - x_O(\tau_O - W_O). \tag{3.16}$$

The first term in both expressions is the change in value due to a gain or loss in bargaining power. By inspection, it is clear that one party gains and one loses. The second term is negative for both parties and represents the loss of otherwise feasible positive NPV projects. So, one party unequivocally loses value, but the other may gain based on the underlying parameters. From (3.15), firm F is net positive with hurdle rate buffers if

$$\frac{\overline{V}}{4} - \theta_N \overline{V} + x_F W_F > 0 \qquad \Longleftrightarrow \qquad \frac{\overline{V}}{4\theta_N} - (\overline{V} - \underline{V}) > 0.$$

This is more likely to be the case if the firm starts with low bargaining power when IRR buffers are not used (feasible). The corresponding condition for the outsider is

$$\frac{\overline{V}}{4} - (1 - \theta_N)\overline{V} + x_O W_O > 0 \qquad \Longleftrightarrow \qquad \frac{\overline{V}}{4(1 - \theta_N)} - (\overline{V} - \underline{V}) > 0.$$

We conclude this section by considering that n parties take part in the investment, and each may act strategically. The primary firm is indexed as $i = 1$, and remainder are in $\{2, \dots, n\}$. For tractability, suppose that the returns (instead of project values) from potential projects are uniformly distributed across the support $R \sim U[0, \overline{R}]$, and that each party's share of the project is

$$\theta_i = \frac{H_i}{\sum_{j \in N} H_j}, \tag{3.17}$$

where $N \equiv \{1, \dots, n\}$. As such, θ_i still takes the form of a Tullock contest function, but we operate in return space as opposed to project value space.

Proposition 3. *For bargaining as in (3.17) there exists a unique symmetric Nash equilibrium in which the optimal buffered hurdle rates are given by*

$$\tau^* = \frac{(n-1)\overline{R}}{n^2} \tag{3.18}$$

and the value for each counterparty is

$$V^* = \frac{\overline{R}}{n^2}. \tag{3.19}$$

The structure of the equilibrium in Proposition 3 resembles that of Proposition 2, except that we now have a comparative static on the number of counterparties at the bargaining table: the equilibrium IRR buffer is strictly decreasing in n.

3.3.3 Non-Strategic Use of IRR Buffers

In this section, we show that the model's main results continue to hold when the model considers non-strategic reasons as the source of IRR buffers. For ease of exposition, in what follows, we continue to assume that the returns to potential projects are uniformly distributed across the support $R \sim U[0, \overline{R}]$. Because buffer use is non-strategic in this section, we assume that the primary firm may use a buffer, but their trading partner does not (i.e., $H_O = W_O$).

Imprecise Estimation of WACC

According to Bessembinder and Décaire (2021) and Krüger et al. (2015), managers may face the inability to precisely estimate W_F. This may be the result of model uncertainty or inherent estimation error. To account for this, or to be conservative, CFOs may use a buffer, which we denote by b.

With a buffer $H_F = W_F + b$, the Nash Bargaining split is

$$\theta_B = \frac{H_F}{H_F + W_O} \tag{3.20}$$

and the value of the firm can be calculated as

$$V_B = \frac{W_F\overline{R} + b\overline{R}}{W_F + b + W_O} - (W_F + b). \tag{3.21}$$

Proposition 4. *When the CFO uses a buffered hurdle rate, $H_F = W_F + b$ to adjust for imprecision in estimating W_F, the value of the firm increases if*

$$V_B > V_N \quad \Longleftrightarrow \quad \frac{\overline{R}W_O - (W_F + W_O)^2}{(W_F + W_O)} > b. \tag{3.22}$$

According to Proposition 4, if the buffer b is sufficiently low, firm value rises because the bargaining benefit outweighs the loss incurred by discarding marginal positive NPV projects. It is instructive to write the inequality in (3.22) as

$$V_B > V_N \quad \Longleftrightarrow \quad (\overline{R} - \underline{R}^N) - \theta_N\overline{R} > b, \tag{3.23}$$

where $\underline{R}^N$ is the lower bound of feasible projects when no IRR buffer is used. This expression indicates that, because the range of projects $(\overline{R} - \underline{R}^N)$ can be considered a measure of uncertainty, according to (3.23) firm value rises when the CFO faces more risk and uses a reasonable buffer to account for lack of precision in her estimate of W_F. Last, according

to (3.23), if the firm already enjoys sufficient bargaining power without a buffer (θ_N is large), it is less likely that using a buffer adds value. In this case, foregoing positive NPV projects is too costly.

Financial and Managerial Constraints

Now, we consider that the CFO chooses a hurdle rate $H_F > W_F$ to accommodate an exogenous constraint that is based on time or financial resources. Suppose that the firm only has financial resources to initiate the fraction f of the potential projects such that $0 < f \leq 1$. Likewise, it can only consider a fraction m of the projects due to managerial constraints such that $0 < m \leq 1$. Define $k \equiv min\{f, m\}$. Given this, the value of the firm with a buffer is

$$V_B = \int_{\underline{R}^H}^{\overline{R}} \theta_H \, dR,$$

where the share of the surplus the firm receives is

$$\theta_H = \frac{H_F}{H_F + W_O}.$$

Proposition 5. *When only k fraction of projects can be considered because of constraints, the CFO sets a buffered hurdle rate*

$$H_F = (1 - k)(\overline{R} - W_O) + k W_F. \tag{3.24}$$

It follows that $H_F > \tau_F^*$ iff

$$k < \frac{\overline{R} - \underline{R}'}{\overline{R} - \underline{R}}, \tag{3.25}$$

where $\underline{R}'$ and $\underline{R}$ are defined as in Section 3.3.2. Finally,

$$V_B > V_N \quad \Longleftrightarrow \quad (\theta_H - \theta_N) > (1 - k)(1 - \frac{W_F + W_O}{\overline{R}}). \tag{3.26}$$

The results in Proposition 5 are straightforward. First, the hurdle rate is a decreasing function of k. By construction, the less constrained the firm is (higher k), the less of an IRR buffer it needs to use. Second, (3.25) tells us that if the firm is sufficiently constrained, it will set a hurdle rate above what is optimal under the strategic considerations in Section 3.3.2, which is suboptimal. However, the use of the IRR buffer due to constraints increases firm value if (3.26) holds. This will arise when the gain in bargaining power is higher, the firm is less constrained, and the range of projects is larger. The latter finding is similar to the risk/variance result that we derived in Proposition 4.

3.4 Model Predictions and Empirical Evidence

In Section 3.4.1, we present empirical predictions from the models described in Section 3.3 and explain how we test these predictions in the data. We then analyze the determinants of the buffer and test model predictions in Section 3.4.2.

3.4.1 Hypothesis Development

We first outline several empirical predictions implied by Propositions 1-5. These predictions are summarized in Table 3.2. Then, we introduce the empirical variables we use to test these predictions while stating each hypothesis. Details on the construction of these variables can be found in Table B.1 and summary statistics in Table 3.1, Panel B.

We start by outlining predictions relating the buffer to four underlying characteristics of firms and projects: the cost of capital, bargaining power, the company's industry, and volatility of project returns. We explore predictions related to both the use of a buffer (hereafter, the extensive margin) and the size of the buffer conditional on use (the intensive margin). A given company or project characteristic (e.g., the cost of capital) is predicted to relate to the extensive margin if it impacts the likelihood the buffer is value-increasing in our model, and is predicted to impact the intensive margin if it impacts the size of the buffer used. Note that for the cases in which the buffer arises exogenously

(Propositions 4 and 5), the model makes no predictions on the intensive margin since the buffer is set for other reasons. Our extensive margin predictions in these cases are based on the exogenous buffer's impact on value.

Hypothesis 1a. *Buffer use is less likely for firms with a higher cost of capital.*

Hypothesis 1b. *Conditional on use, the size of the buffer is smaller for firms with a higher cost of capital.*

Note that Hypothesis 1 distinguishes our theory from other buffer explanations because the predictions differ from what we would expect for explanations that do not consider the impact of the cost of capital on bargaining. For example, if managers use a higher buffer to compensate for idiosyncratic risk (Décaire, 2021), we would expect there to be, on average, a positive association between the cost of capital and the buffer – because empirically idiosyncratic risk is positively correlated with systematic risk. Prior explanations based on financial or managerial constraints would also generally predict either no relation (since managerial time constraints are likely independent of the cost of capital) or a positive relation (since financial constraints are more likely to bind for high cost of capital firms).

In contrast, Hypothesis 1 indicates that our model predicts a negative relation between the hurdle rate buffer and the cost of capital because firms with a higher cost of capital already have more bargaining power, all else equal. In this situation, any incremental benefit of more bargaining power would be outweighed by the value lost from good projects ruled out by a higher buffer, on net leading to a negative relation between the cost of capital and the buffer.

A second hypothesis relates to our model predicting a systematic relation between the buffer and a firm's bargaining power. For the single firm case, we can use (3.10) to predict that increasing bargaining power decreases use of a buffer. For the intensive

margin, according to Proposition 1,

$$b = \frac{\sqrt{\overline{V} x_O W_O} - x_O W_O}{x_F} - W_F,$$

and therefore $\frac{\partial b}{\partial W_F} < 0$. If we hold, x_O, x_F, and W_O fixed, then $\frac{\partial W_F}{\partial \theta_N} > 0$. This implies that $\frac{\partial b}{\partial \theta_N} < 0$. So, because θ_N is a measure of ex-ante bargaining power, we expect the buffer to decrease in bargaining power. In the competitive model, according to (3.13), ΔV is decreasing in θ_N, so on the extensive margin, we expect firms to use buffers less. Also, according to Propositions 4 and 5, in the versions of the model where managers set hurdle rate buffers for non-strategic reasons, higher bargaining power is associated with decreased likelihood that a buffer will create value through a bargaining channel: The left-hand side of (3.23) is decreasing in θ_N, as is equation (3.26).

Hypothesis 2a. *Buffer use is less likely when a firm has more bargaining power.*

Hypothesis 2b. *Conditional on use, the size of the buffer is smaller for a firm with higher bargaining power.*

These predictions are not in the traditional buffer literature as there is no immediate analogous prediction related to bargaining power from buffer explanations based on idiosyncratic risk, financial constraints, or managerial constraints. To test Hypothesis 2, we measure bargaining power in two ways. First, using financial disclosures on corporate customers, we measure the concentration of a firm's customer base. The higher the customer concentration, the more reliant the firm is on a single customer, and therefore the lower is the firm's bargaining power, all else equal. Second and separately, we measure markups using the methodologies presented in Baqaee and Farhi (2020b) and De Loecker et al. (2020), where higher markups proxy for more bargaining power. In our analysis of markups, which we explain in detail below, we focus on firms whose transactions are predominantly business-to-business (B2B) because the price markup of B2B firms is a clear consequence of bargaining power over customers/suppliers.

Our third hypothesis comes from the basic premise of the model that project value is a function of bargaining over inputs and outputs. We expect that a buffer is less likely to arise or increase value in, for example, industries where inputs are purchased at liquid market prices. Therefore, we predict that buffer use will be positively correlated with industries where a given project's inputs or outputs are large and their value is specific to the project (e.g., land).

Hypothesis 3a. *Companies are more likely to use a hurdle rate buffer if they operate in industries where projects require more property, plant, and equipment (PPE).*

Hypothesis 3b. *Conditional on use, the size of the buffer is larger in industries where more PPE is required for projects.*

We note that firms bargain over more than just tangible assets. Our framework readily applies to, e.g., a manufacturing firm purchasing supply chain management software. Hypothesis 3 is an empirical prediction under the view that bargaining plays a more important role when the tangibility of assets is higher.

Our predictions in Hypothesis 3 are new and generally differ from predictions implicit in non-strategic buffer explanations (that do not account for the effect of bargaining on project outcomes), thus distinguishing our analysis from these other buffer explanations. For example, because financial constraints are decreasing in tangibility (Almeida and Campello, 2007), explanations for the buffer based on such constraints alone would lead to the opposite prediction. We similarly would expect the opposite prediction if the buffer is a response only to idiosyncratic risk, since idiosyncratic volatility is expected to decrease in tangibility (Fink et al., 2010). Since there is no obvious relation between managerial time constraints and tangibility, the time-constraint explanation for the buffer has no clear prediction on this dimension absent the bargaining effect.

To test Hypothesis 3, we use two proxies to identify industries in which bargaining is likely to play a role. First, we group firms into industry buckets based on the type

of underlying investments (e.g., fixed assets vs human capital). The four groups are: 1) Manufacturing, Construction, Mining, Energy, 2) Finance, 3) Services, Communications, Media, and 4) Healthcare, Pharma, Tech. Second, we measure each industry's average asset tangibility (defined as net PPE scaled by total assets) of Compustat firms at the NAICS-4 industry level.[20]

In addition to this methodology, we also explore this hypothesis with respect to the redeployability of assets (Kim and Kung, 2017). All else equal, we expect that assets with lower redeployability (more specific-to-use) are more likely to be bargained over, which increases the incentive to use a buffer in negotiations.

Our last prediction is that buffer use and size generally increase with the volatility of project returns. According to Propositions 4 and 5, ΔV is increasing in $\overline{R} - \underline{R}^N$. The predictions from the single firm strategic case and the competitive buffer case (Proposition 2) are a little more nuanced. Because the lower bound on acceptable projects is fixed, an increase in $\overline{V}$ leads to both a higher average project value and increased variance. Clearly, τ^* is increasing in $\overline{V}$ in both strategic models. So, we do predict that project volatility is positively correlated with buffer use, keeping in mind the limitations of those versions of the model. Notwithstanding, previous work predicts this as well. Décaire (2021) argues that firms use buffers to compensate for idiosyncratic volatility. Likewise, buffers that are put in place due to financial constraints may be because a firm faces uncertainty.

Hypothesis 4a. *Buffer use is more likely for firms with higher volatility of underlying project values.*

Hypothesis 4b. *Conditional on use, the size of the buffer is larger for firms with volatility of underlying project values.*

To measure volatility, we adapt the methodology used in Décaire (2021). Specifically,

[20] A regression of firm-level tangibility on NAICS-4 fixed effects gives an R^2 of 0.62, whereas including year fixed effects only increases this to 0.635, thus there is little gained from using a time-varying industry-level measure of tangibility.

we compute firm-level idiosyncratic volatility as the residual from a regression of sales to lagged assets on NAICS-4 and quarter fixed effects. We then measure industry volatility as the standard deviation of these residuals within a CFO survey firm's NAICS-4 industry the quarter previous to the survey quarter. This measure captures the dispersion of firm-level idiosyncratic volatility within a CFO survey firm's industry. If managers alter their discount rates in response to volatility, then this measure should correlate positively with the buffer.

3.4.2 Empirical Analysis of Model Predictions about the Hurdle Rate Buffer

We now provide deeper analyses of the hurdle rate buffer, guided by the empirical evidence from Section 3.2 and the model predictions made in Section 3.4.1. In our regression analyses, all continuous variables are standardized to unit variance except for the firm's cost of capital. Table 3.1, Panel A displays summary statistics for variables related to the buffer, and Panel B displays the same for other variables; Figure B.1 presents pair-wise correlations.

In Table 3.3, we analyze the determinants of the whether firms use a hurdle rate buffer and if so, how large it is. Panel A presents analysis concerning the extensive margin of whether a firm uses a buffer; Panel B displays evidence about the factors that drive the intensive margin (the magnitude of the buffer). Because the specifications in both panels are similar, we discuss them together.

Hypothesis 1. *The Firm's Cost of Capital*

As can be seen in Table 3.3, both the probability of using a buffer and the size of the buffer itself are strongly negatively related to the cost of capital.[21] These negative relations support Hypothesis 1. A one percentage point increase in the cost of capital predicts a 2.5 to 3 percentage point (pp) drop in the probability of using the buffer, and

[21] These negative relations are very robust, and exists in every specification in the paper.

roughly a 0.3 pp reduction in size of the buffer.[22]

This negative correlation between cost of capital and buffer use is noteworthy with respect to several existing explanations of why firms set hurdle rates above WACC (as presented in Figure 3.1). According to these explanations, if either uncertainty about a firm's true WACC or project-specific idiosyncratic risk are large, then we would expect the cost of capital and the buffer to be *positively* correlated: As the required return of projects increases, so too would these two sources of uncertainty; therefore, managers would build in a *larger* buffer as the cost of capital increases.

Based on estimates from Table 3.3, Figure 3.6 illustrates the negative relation between the cost of capital and the buffer. Panel A displays predicted probabilities that a firm uses a positive buffer for different costs of capital. Going from a cost of capital of 5% to 15%, the probability that a company uses a buffer drops from about 87% to 60%. Panel B displays a binned scatter plot of the size of a firm's (intensive margin) buffer and its cost of capital. A strong negative relation between the two variables is apparent. We note that the relations shown in the figure are net of both size and industry fixed effects (the underlying regressions control for both). While the model makes predictions about the size and likelihood of the buffer across industries, the negative relation between the buffer and the cost of capital should exist for all types of firms that strategically set buffers.

A natural question is whether the negative relation between the buffer and the cost of capital holds within the subset of firms with high cost of capital. Because the number of projects decreases in the project's return, it may be that high WACC firms are forced to decrease their buffers to find feasible projects. This could lead to less variation in the buffer among the firms with higher costs of capital, as the buffer gets forced down to zero.[23] Figure B.4 estimates the relation between the buffer and WACC for high, medium

[22] If we condition on firms that set their hurdle rate above WACC, the effect of a one pp increase in the cost of capital leads to about a 0.18 pp drop in the buffer, as in columns 5-8 of Table 3.3, Panel B.

[23] Variation might also be limited if project prioritization were the primary driver of the buffer. Low WACC firms can easily prioritize projects, whereas high WACC firms cannot.

and low WACC firms and shows that the negative relation exists within each subset, even among firms with a high cost of capital (15% or greater).

Control Variables in Table 3.3 Other research shows that estimation error or uncertainty about WACC and/or discount rates affects corporate investment (Krüger et al., 2015; Bessembinder and Décaire, 2021). While not directly related to our bargaining model, this would affect hurdle rate buffers in our data, if CFOs include a fudge factor in their discount rates to offset this uncertainty. We use the volatility of the CAPM beta within the CFO survey firm's NAICS-4 industry the quarter previous to the survey quarter to proxy for this uncertainty (Bessembinder and Décaire, 2021).[24] Table 3.3 shows that beta volatility is positively correlated with both the use and magnitude of the buffer.

Our proxy for idiosyncratic project risk, industry sales volatility, is strongly related to the intensive margin of the buffer, but is not significant on the extensive margin (see Hypothesis 4). Combined with the previous paragraph, this suggests that, after controlling for estimation uncertainty in WACC, even for firms that estimate their cost of capital very well, idiosyncratic project risk leads to larger buffers.

Large firms are both less likely to use a positive buffer and have smaller-sized buffers, reinforcing the findings from Figure 3.5, Panel C. Public firms are slightly more likely to use a positive buffer, but show no difference in the size of the buffer. Whether or not a firm has a credit rating has no effect on either margin. To the extent that having a credit rating is a reasonable measure of financing constraint (Faulkender and Petersen, 2005), we do not find evidence in our sample that the ability to access public debt markets is a primary determinant of the buffer.

The CFO survey asks CFOs to rate their level of optimism about their own firm's

[24] When this measure is high, firms within an industry vary substantially in the level of return investors require in order to bear the firm's underlying risk. Even for firms not publicly traded, this measure proxies for the required rate of return demanded by investors if the business activities and underlying cash flow risk of private and public firms within an industry are similar.

prospects and the US economy more generally.[25] Columns 4 and 7 of Panel A and columns 4 and 8 of Panel B include these variables in our baseline regressions. We find that own-firm optimism is positively related to buffer use and the size of the buffer, even after controlling for US economy optimism. This suggests that CFOs choose a larger buffer at least in part to offset over-optimism in other aspects of project development. We note that in contrast, if firm-optimism is correlated with higher confidence of project success and this were the dominant effect, we would expect to find a negative relation between optimism and the buffer.

The Buffer, the Cost of Capital, and Beta Table 3.3 shows that that the use and size of the buffer are strongly negatively related to the cost of capital. This is a key prediction of our model – all else equal, two firms with different costs of capital will have different incentives to use a buffer. Implicit in this prediction is the assumption that two firms can face the same distribution of project returns, but have different costs of capital. This assumption is consistent with a CAPM determined cost of capital as long as not all project volatility is systematic. Nonetheless, it is important empirically to ensure our results are not driven by systematic differences in the risk of projects.

With this in mind, we conduct additional tests of the relation between the buffer and WACC, controlling for firms' estimated betas. To do this, we rely on the publicly traded firms that appear in the CFO survey; for which we can directly estimate their CAPM betas (see Table 3.4).The analysis shows that the negative buffer-WACC relationship continues to hold, as before. The table also provides further reassurance that our survey data are sound – survey-provided discount rates are positively related to the firm's beta (consistent with Jagannathan et al., 2016).

Hypothesis 2. *Bargaining Power in Negotiations*

[25] The question asks CFOs to rate their level of optimism on a 0-100 scale, with 0 being the least optimistic. The mean firm-level (US) optimism is 68.7 (63.1), with a standard deviation of about 18 (16).

Our second main hypothesis predicts a systematic negative relation between the size and use of the buffer and a firm's degree of bargaining power when negotiating with customers and suppliers. Bargaining power in the model is relative, thus in our empirical analysis we analyze the firm's *own* level of bargaining power, and separately that of its trading partners.

Customer Concentration As a first step, we directly quantify the relative bargaining power of a given company by measuring it relative to the market position of its customers. We proxy for the bargaining power of a seller/supplier based on the level of concentration of the firm's sales. We use data from the Compustat Segments file to measure firm-level customer sales concentration, which pulls data from company filings about firms' large customers. Specifically, we calculate a Herfindahl–Hirschman Index (HHI) of each firm's sales to its corporate customers (Patatoukas, 2012). To match to our survey data, we take quarter × NAICS-3 averages of this measure and match to each firm in the survey by its industry.[26]

The measure is bounded between zero and one – when it is closer to zero, the sales of firms in an industry are spread more evenly across a larger number of corporate customers, indicating more bargaining power for the firm *relative to* its customers. Likewise, given that this is a *customer*-based measure of concentration, when this measure approaches one, the seller firm has *less* bargaining power, all else equal, because the seller is more reliant on fewer customers to generate revenue.

In Table 3.5, we relate this measure to survey firms' buffers. Columns 1-3 analyze the extensive margin; columns 4-6 analyze the intensive margin. We find no relation between customer concentration and the extensive margin of the buffer. In terms of the

[26] The industry level for most of the variables in our analysis is NAICS-4; however, our bargaining power measures are at the NAICS-3 level. We do this to align with the industry level that Baqaee and Farhi (2020b) focus on in their analysis, and because we estimate the output elasticities used in Table B.5 Panel C at the NAICS-3 level. Table B.1 provides detail on the exact construction.

intensive margin, columns 4-6 show that a one-standard deviation increase in customer concentration leads to about a 0.3 pp increase in the buffer, consistent with the model's prediction. For context, our results predict that moving from the 10th to 90th percentile of customer concentration would increase the size of the buffer by about 0.9 pp.[27]

Markups Next, we proxy for the firm's own bargaining power by using markups. We note that markups in B2B transactions are a consequence of bargaining between trading partners (either customers or suppliers) and thus are close in spirit to the model. The model predicts a negative relation between the buffer and bargaining power as proxied by B2B markups.[28]

At consumer-facing firms (in contrast to B2B), price markups are over consumers, and may be more determined by competition in the firm's local market (and less reflective of the model's bargaining power over other firms). Therefore, we would not necessarily expect the model's predicted negative relation between markups and buffers to hold in that setting. In our main analysis of the relation between the buffer and markups, we therefore study a subsample of firms in B2B industries to better isolate the prediction in Hypothesis 2. Following Gofman et al. (2020), we classify industries as consumer-facing if they are in GICS sectors "Consumer Discretionary" or "Consumer Staples." In the CFO survey data, 13.5% of observations are in consumer-facing industries, which we delete from the B2B subsample analysis.

Our primary markup measure is the "accounting" markup developed in Baqaee and

[27] We use industry-level measures because the CFO survey sample is comprised of many private firms, so we cannot estimate their customer concentration directly. To justify this approach, we do two tests. First, we regress Compustat firm-level customer concentration on its industry-by-time average counterpart, yielding an R^2 of 0.41, suggesting customer concentration has an important industry component. Second, using the CFO survey firms in Compustat (15.1% of our sample), we do similar analysis to Table 3.5, substituting the industry customer concentration for the firm's customer concentration; coefficients are directionally the same with similar magnitudes, though the statistical significance is not as uniformly strong, perhaps due to lower power in the smaller sample. Results are available in Table B.6.

[28] Markups are commonly used as an empirical proxy for market power, on which our model makes no prediction. The price markup for B2B firms is an adequate, if noisy, proxy of bargaining power. Nevertheless, we cannot perfectly separate these two forces (bargaining and market power).

Farhi (2020b) and we supplement this analysis using the methodology in De Loecker et al. (2020). To leverage our full survey sample, based on Compustat firms we aggregate markups to the NAICS-3 by year and match to our survey firms.[29]

Table 3.6 displays the markup-buffer analysis related to Hypothesis 2. The table shows that a one standard deviation increase in markups reduces the likelihood of using a buffer by 2 pp and decreases the size of the buffer by about 0.5 pp. These results predict that moving from the 10th to the 90th percentile markup would reduce the size of the buffer by 1.1 pp.[30] In Table B.5, Panel C, we perform similar analyses where we estimate markups via the methodology in De Loecker et al. (2020), still focusing on firms in B2B industries. The results are similar.

As discussed above, the model makes no prediction about the relation between the markup and the buffer for firms in consumer-facing industries. But for completeness, in Table B.5 Panel A, we repeat the analysis of Table 3.6 using the full CFO Survey sample. While results are similar, the coefficients on markup are slightly lower (towards zero) by including consumer-facing firms. In Table B.5 Panel B, we explore this further and interact the markup variable with a consumer-facing indicator to test if markups affect buffers differently in consumer-centered industries. The table displays group-specific slopes and shows for non-consumer-facing (i.e., B2B, as in our main analysis), the signs for extensive and intensive margins are negative, as expected.

Hypothesis 3. *Industry Differences and Asset Tangibility*

We document the use and size of the buffer across general industry classifications, to test the basic premise of the model. Table 3.7 Panel A shows that firms in manufacturing, construction, mining and energy industries are 7.3 percentage points more likely to

[29] A regression of firm-level markup on its industry-by-time counterpart yields an R^2 of 0.46, suggesting that price markups have an important industry component.

[30] Table B.6 shows that coefficients are of similar magnitude when we do similar analysis using firm-level markups for CFO Survey firms that appear in Compustat.

use a buffer, and also use significantly larger buffers (0.79 pp larger) than firms in other industries, even after controlling for industry risk and the firm's cost of capital. These are the types of industries in which project inputs tend to be larger (e.g., investment in machinery) and project-specific, consistent with the spirit of the model.

We complement this analysis by examining the relation between the buffer and industry tangibility in Table 3.7 Panel B. This provides a direct measure of industry-level project inputs to strengthen the analysis in Panel A. Panel B shows that a one standard deviation increase in asset tangibility (measured at the industry level) increases the probability of using a buffer by approximately 3.5 pp and increases the size of the buffer by 0.3-0.4 pp.

While many tangible assets are naturally specific-to-use, with prices more likely to be influenced by the relative bargaining power of trading partners, other forces can play a role in the documented positive relation. With this in mind, we refine our analysis to also explore *asset redeployability*, as in Kim and Kung (2017), which captures the usability of assets within and across industries – high redeployability implies an active secondary market and hence more competitive prices (and less room for a bargaining advantage). We thus expect firms that purchase and use assets with *lower* redeployability have a buffer, as these assets are more specific to utilize.

We match the industry-level measure from Kim and Kung (2017) based on 2-digit NAICS codes; Table B.7 displays the analyses. We find a negative relation on both margins, with the intensive margin being more robustly significant – firms that rely on assets with lower redeployability have larger buffers, all else equal (the negative coefficient estimate for redeployability makes sense, given the variable's strong negative correlation with tangibility in our data, as expected). Though not shown in the table, the statistical significance of the relation between redeployability and the buffer is weaker at finer industry codes (e.g. NAICS-3), possibly due to measurement error from using finer classifications.

Hypothesis 4. *Volatility of Underlying Project Values*

As explained in Section 3.4.1, the model predicts a higher likelihood of using a positive buffer and a larger buffer if a firm's underlying project risk is higher. We refer to Table 3.3 to discuss the role of this project risk (proxied by sales volatility), and we also control for this variable in many specifications. While our model predicts a positive correlation between project volatility and the extensive margin of the buffer, we find no relation in Panel A. However, industry sales volatility has a large effect on the intensive margin of the buffer – a one standard deviation increase in sales volatility leads to a $\approx$ 0.4 pp increase in the size of the buffer (Panel B), consistent with the model and Décaire (2021). Further, the positive estimated coefficient for beta volatility is also consistent with a positive relation between idiosyncratic project (or cash flow) volatility and the buffer, given the positive relation between beta and idiosyncratic cash flow volatility documented in Babenko et al. (2016).

3.5 Conclusion

Project development differs from standard textbook capital budgeting guidance. We explore these differences based on two observations about real-world capital allocation. First, input costs are not exogenous as commonly assumed, but rather are the result of negotiation between a firm and its trading partners. Second, headquarters tells mid-level managers the hurdle rate they must use for project evaluation, which the managers accept without question. As is well known, and confirmed by our survey data over the past dozen years, CFOs typically inflate the hurdle rate with a buffer of 5 or 6 percentage points above the cost of capital. We also provide direct empirical evidence that hurdle rates are a focal point of corporate investment decisions and that companies exhibit "beat the benchmark" behavior relative to the hurdle rate.

We build a model to explore the implications of these realistic features of project development. The model highlights a trade-off that occurs from using an inflated hurdle

rate: passing up projects with moderately positive net present value versus earning a higher share in negotiations from the smaller set of projects. We explore these issues in several settings, including where firms set inflated hurdle rates (i) for non-bargaining reasons, (ii) strategically to gain bargaining power, and (iii) where firms and their trading partners both use hurdle rates for bargaining. In the model, the inflated hurdle rates reflect firms' relative bargaining power, their relative investment in the project, and the value of the underlying projects.

In equilibrium, using an inflated hurdle rate can preserve or even create value, thus offering an explanation for the wide-spread and long-standing use of inflated hurdle rates. Importantly, the implications of our model hold whether companies strategically choose inflated hurdle rates to aid negotiation or whether hurdle rates are inflated for other "non-strategic" reasons posited by the literature. Moreover, even when a counterparty is using an inflated hurdle rate for strategic reasons, a firm will find it optimal to use a hurdle rate buffer.

We test four predictions of the model using data that span 2011 to 2022 provided by the expert practitioners who make project development decisions. Our empirical findings generally support the model's predictions. We find that hurdle rate buffers are negatively correlated with the cost of capital and negatively correlated with the firm's ex ante relative bargaining power. These two predictions are novel and opposite from what one would expect based on existing hurdle rate buffer models. We also show that hurdle rate buffers are more likely when firms (or industries) rely on fixed assets and find evidence of a positive relation between volatility of sales and the hurdle rate buffer as predicted. Our empirical exploration of these latter two predictions based on CFO survey data is new, though the predictions themselves align with other explanations of the hurdle rate buffer.

Future research should explore other realistic features of project development to determine whether and how these features affect project choice and valuation. Consider, for

example, several real-world features documented in Graham (2022). First, hurdle rates are typically sticky through time, not moving much as market conditions change. Our paper has explored how the hurdle rate buffer affects project development; future research can examine the effects of hurdle rate stickiness. Second, the budgeting process is discrete, lumpy and rigid. Décaire and Sosyura (2022) explore the impact of budgeting rigidity, showing that as the year-end approaches, companies that have underspent to date on advertising expenses typically go on a year-end spending spree that leads to inefficient investment. Third, the horizon over which companies can realistically plan and forecast cash flows has shrunk in recent years and project life has accordingly gotten shorter. Dessaint et al. (2022) confirm this result using textual analysis of financial filings. Relatedly, managers often require that a project pay off within two or three years, similarly affecting project development via a short-term focus. In addition to a short planning horizon, Graham (2022) shows that internal forecasts for important corporate planning variables are remarkably inaccurate. Finally, corporate managers are typically miscalibrated (underestimate both the upside and downside likelihoods) and also use simple decision rules.

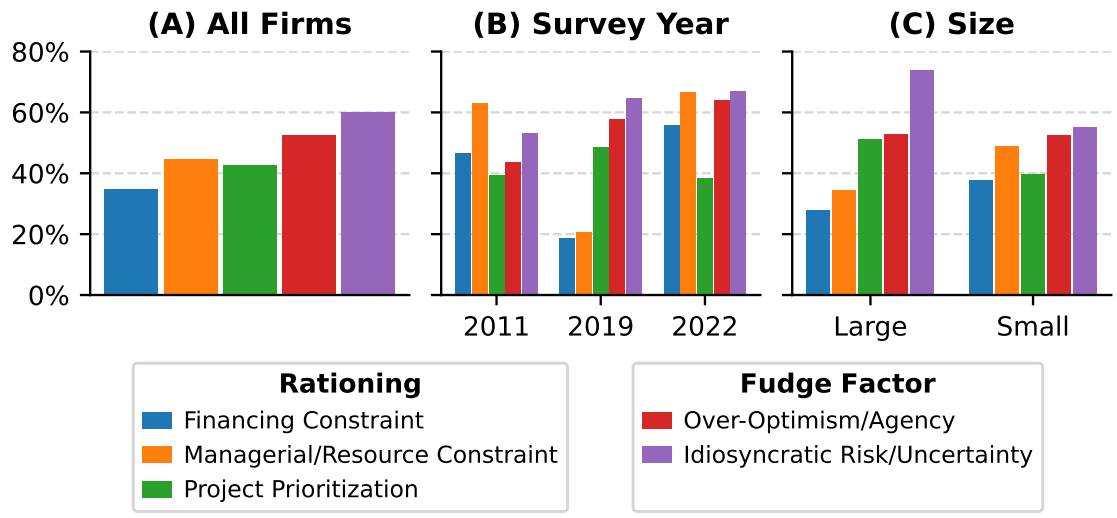

FIGURE 3.1: Reasons that Companies Set Hurdle Rates Above the Cost of Capital

This figure displays qualitative reasons why firms set their hurdle rates above WACC (i.e., use a hurdle rate buffer). Results are taken from three editions of the Duke CFO survey (2011q1, 2019q1, 2022q2) which asked CFOs "Why do you set your hurdle rate above WACC?" Possible answers vary by survey, so we group them into five distinct reasons. Table B.2 displays these groupings and provides further detail on how the question was asked on each survey. "Rationing" refers to cases when managers prioritize projects, for example due to limited capital resources. "Fudge Factor" refers to cases when managers add a (positive) buffer to their discount rate to account for, e.g., idiosyncratic risk of specific projects (Brealey and Myers, 1996; Décaire, 2021). Panel A displays the percentage of CFOs that fall within each category, Panels B and C displays the results split by survey year and size (revenue above/below $1 billion).

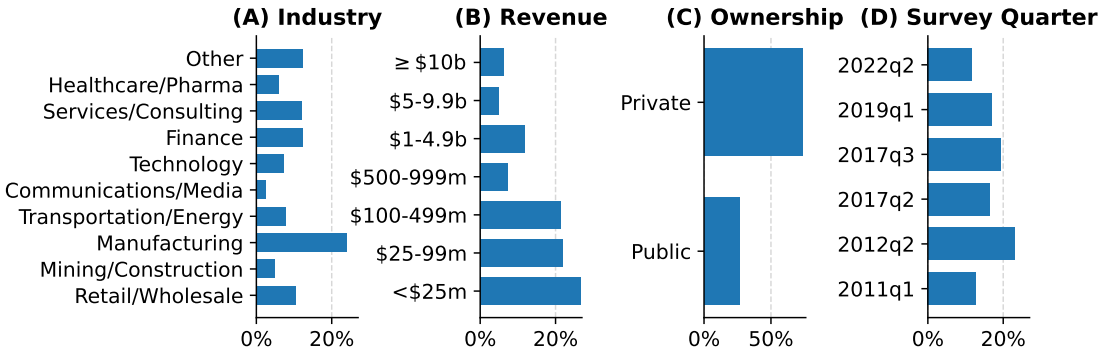

FIGURE 3.2: Sample Demographics

This figure displays demographic information for all CFOs that provide both a hurdle rate and cost of capital (WACC) in the Duke CFO Survey.

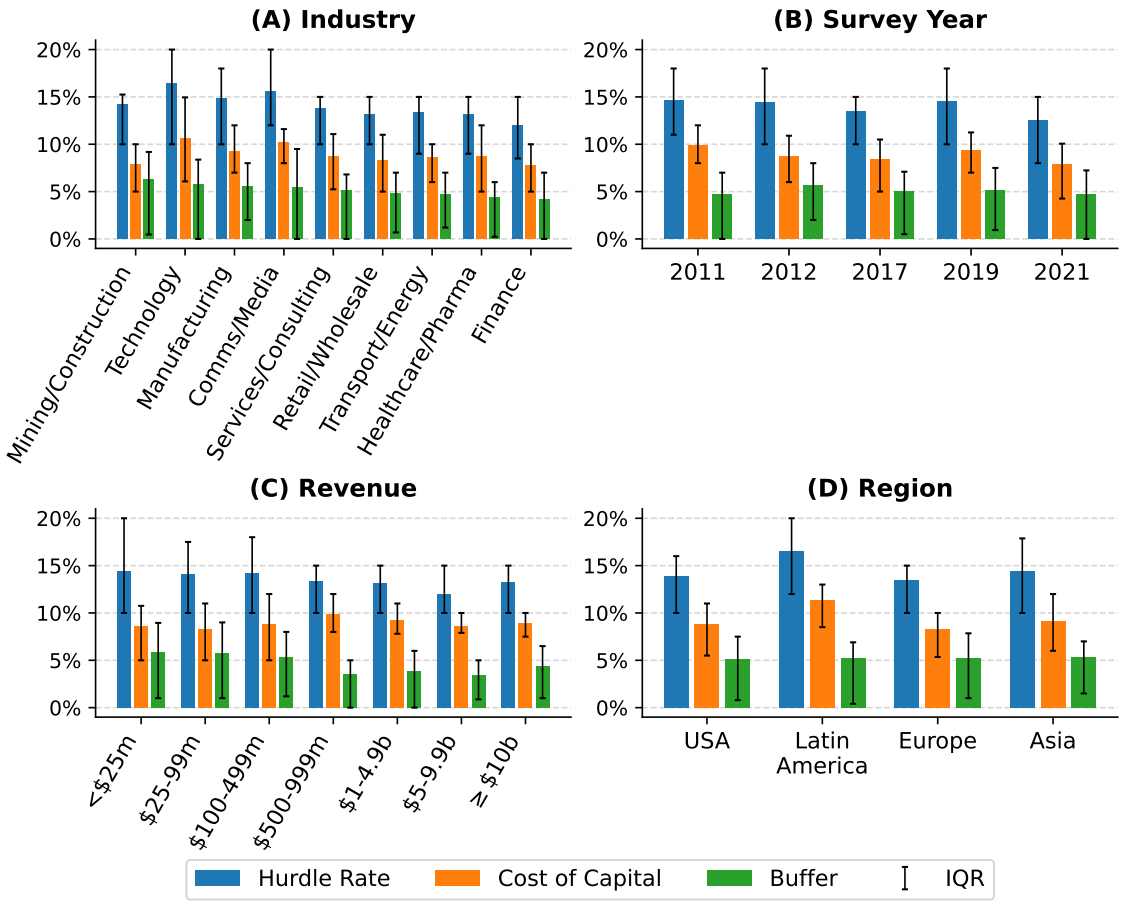

FIGURE 3.3: Time Series and Cross-Sectional Variation in Hurdle Rates and the Cost of Capital

This figure displays the within-group average and interquartile range (IQR) of CFO survey company hurdle rates, WACCs and hurdle rate buffers across several observable characteristics. For example, Panel A displays the statistics across CFO survey industries; the average buffer for firms in Mining/Construction is about 6%, and the 25th and 75th percentiles within-group are roughly 0% and 9%, respectively.

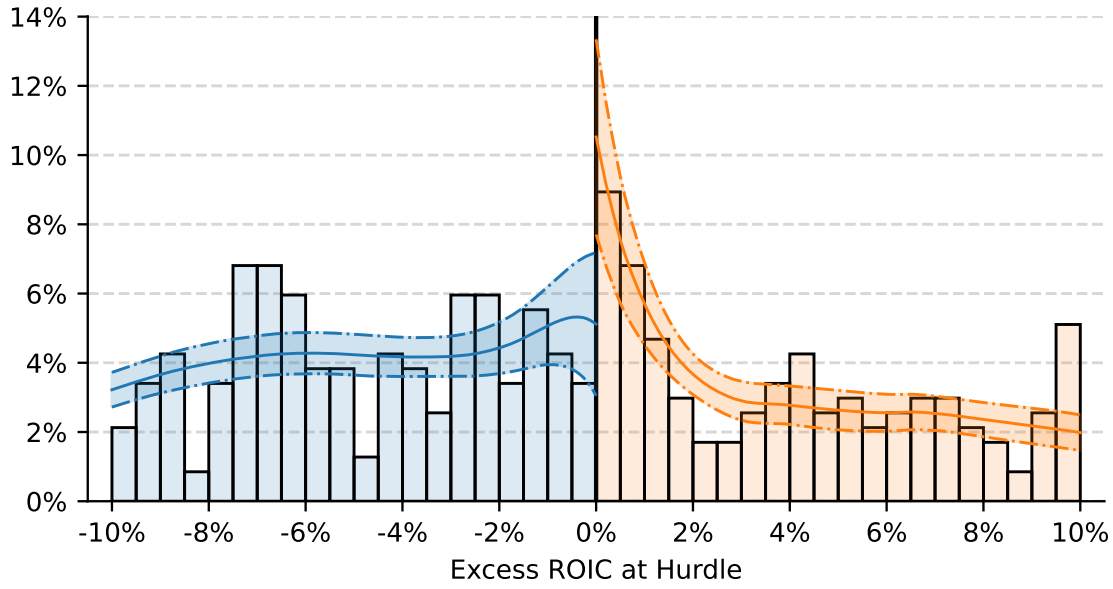

FIGURE 3.4: Density Manipulation in Excess ROIC at the Hurdle Rate

This figure displays the results of testing for density manipulation of return on invested capital (ROIC) around the hurdle rate, following the methodology described in Cattaneo et al. (2018). The sample is CFO survey firms that also appear in Compustat quarterly, for which we observe a hurdle rate. We first linearly interpolate firms' hurdle rates between two consecutive appearances in Compustat. For example, if a firm's hurdle rate in 2012q2 were 10%, and in 2012q4 it were 11%, we would assume that the value of the hurdle rate in 2012q3 were 10.5%. Given stickiness in hurdle rates (as in, e.g., Graham, 2022), this assumption seems reasonable. ROIC is defined as $ROIC_{i,t} = EBIT_{i,t}(1 - tax_{i,t})/ICAPT_{i,t-1}$, where $ICAPT$ is the firm's invested capital (i.e., the sum of long-term book debt and equity). We take trailing four-quarter sums of $EBIT(1 - tax)$, and divide by $ICAPT$ from the previous fiscal year. Next, we define our variable of interest, "Excess ROIC," as a firm's ROIC above its hurdle rate, i.e. for firm i in quarter t,

$$Excess\ ROIC_{i,t} = ROIC_{i,t} - Hurdle_{i,t}$$

We focus on observations of Excess ROIC with an absolute value less than 25%. The sample is 470 observations. We are interested in testing for density manipulation of Excess ROIC at 0. The blue and orange bars display observed frequencies of Excess ROIC in 0.5% bins ranging from −10% to +10% (though the densities estimated use the full range of observations). The blue and orange lines (and shaded areas) display the estimated local-polynomial densities as described in Cattaneo et al. (2018). To estimate the densities, we set the order of the local polynomial (and the order of the bias-corrected density estimator) equal to 3. The bandwidths for the local polynomial estimator are chosen via the "data-driven" methodology described in Section 2.5 of Cattaneo et al. (2018), and are $\{h_-, h_+\} = \{7.449\%, 6.199\%\}$. The shaded areas display 90% confidence intervals. Table B.4 displays robustness where we vary the size of the bandwidth and the order of the polynomials. The figure displays a discontinuity of the ROIC density at the hurdle rate, with some mass of the ROIC density shifting from just below to just above the hurdle rate.

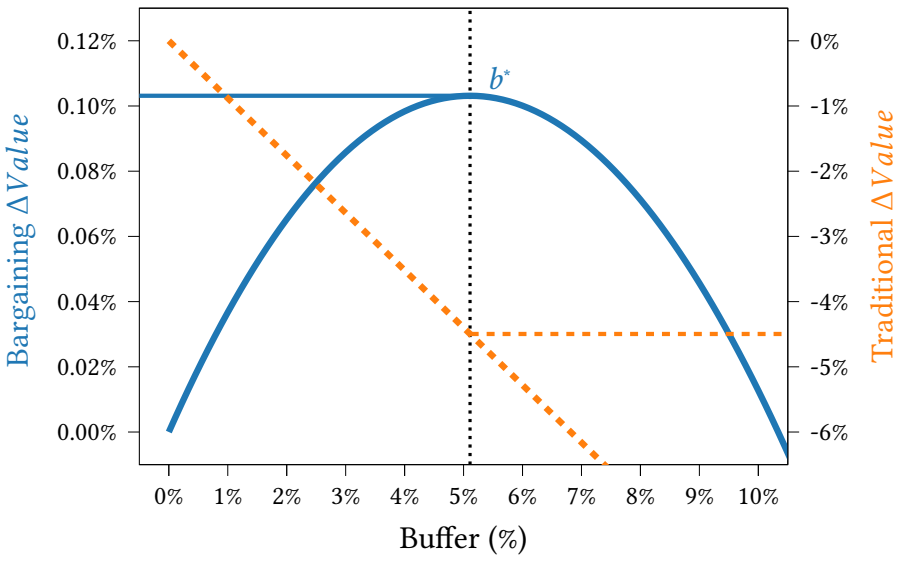

FIGURE 3.5: Implications for Firm Value from the Delegated Bargaining Model

This figure plots the model-implied change in firm value from using a hurdle rate buffer and the change in firm value from using a buffer as implied by a standard capital budgeting model. Using Proposition 1, we calibrate our model such that $r_F^{WACC} = 8.77\%$ and $r_O^{WACC} = 13.88\%$, the mean and 90th percentile WACC in the CFO Survey data. We choose remaining parameters such that the optimal hurdle rate buffer $b^* = 5.11\%$, which is the average buffer in our sample (in the graph, we set $x_F = x_O = 1$ so that firms have equal asset sizes, which gives $\underline{V} = 2.228$; and we set $\overline{V} = 4.556$, about $2 \times \underline{V}$). Under these parameters, the percentage increase in firm value from the no-buffer case to the optimal buffered hurdle rate, $V_B^*/V_N - 1 \approx 0.11\%$, and is displayed in blue relative to the left y-axis. Displayed on the right y-axis is the percentage change in value if we were to ignore the benefits from bargaining (and only consider the value lost from forgoing positive NPV projects), which is roughly −4.4%. Thus, the difference in value implications once we consider the effects of bargaining is about 4.5 percentage points.

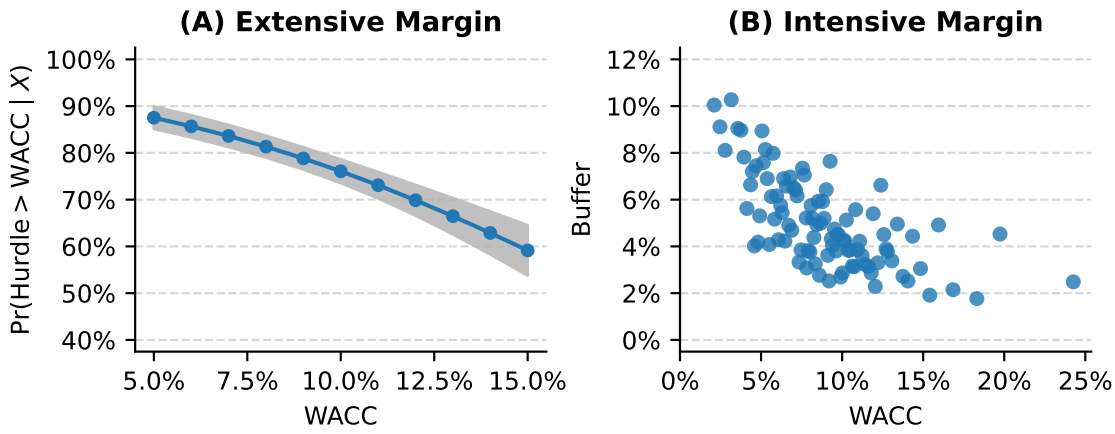

FIGURE 3.6: The Relation Between the Cost of Capital and Hurdle Rate Buffers

This figure displays the effect of a firm's cost of capital on the extensive and intensive margins of the buffer. Panel A displays estimated probabilities that a firm uses a buffer for different costs of capital (estimated using column 6 of Table 3.3, Panel A). Each navy point on the plot displays the estimated probability of a positive buffer, given that level of cost of capital. For example, for a cost of capital of 5%, $\Pr[Hurdle > WACC] \approx 87\%$, and this probability decreases with WACC. The gray shaded areas are 95% confidence intervals. Panel B displays a binned scatter plot with 100 bins of a firm's (intensive margin) buffer and the cost of capital (estimated using column 4 of Table 3.3, Panel B), again showing a negative relation with WACC.

Table 3.1: Summary Statistics

	N	Mean	Std Dev	25%	50%	75%
Panel A: CFO Survey Variables						
Hurdle Rate	1,232	13.879	5.977	10	12.750	16
Cost of Capital	1,232	8.769	4.211	5.500	8.800	11
Buffer (Extensive Margin)	1,232	0.777				
Buffer (Intensive margin)	1,232	5.111	5.376	0.787	4	7.500
Buffer | Buffer > 0	957	6.579	5.248	3	5	9.500
Panel B: Industry-Level Variables						
Beta Volatility	1,232	0.686	0.234	0.535	0.687	0.834
Sales Volatility	1,232	0.116	0.050	0.088	0.117	0.148
Customer Concentration	1,232	0.292	0.198	0.113	0.316	0.426
Markup	1,066	1.091	0.186	1.015	1.058	1.108
Markup (Alternative)	1,066	1.093	0.102	1.055	1.105	1.154
Asset Tangibility	1,232	0.231	0.159	0.120	0.194	0.297
Asset Redeployability	1,232	0.440	0.128	0.346	0.476	0.516

This table displays the number of observations, averages, standard deviations, and quartiles of key empirical variables. Panel A displays statistics for variables related the hurdle rate and the cost of capital. Panel B displays summary statistics for the industry-level variables we explore in Section 3.4. Summary statistics for markup variables are displayed only for firms in non-consumer-facing industries. Detailed variable definitions are available in Table B.1.

Table 3.2: Empirical Predictions of Model Propositions

Hypothesis Buffer Use	1 Cost of Capital		2 Barg. Power		3 Asset Tang.		4 Vol. of Projects	
	Int.	Ext.	Int.	Ext.	Int.	Ext.	Int.	Ext.
Strategic 1-firm	–	–	–	–	+	+	+	+
Strategic 2-firm	–	–		–	+	+	+	+*
Exog. buffer - uncertainty	–	–	–	–	+	+	+	+
Exog. buffer - constraints	+/–	+/–	–	–	+	+	+	+

*: If ΔV is positive, then it must be that $\theta < \frac{1}{2}$.

This table summarizes the direction of the predicted relation between buffer use on the intensive margin (Int.) and extensive margin (Ext.) and four firm characteristics: cost of capital, bargaining power, asset tangibility and volatility of project returns. The predicted sign of a given relation is based on Propositions 1-5 in our model. The columns of the table align with Hypotheses 1, 2, 3, and 4 respectively, as presented in Section 3.4.1.

Table 3.3: Determinants of the Extensive and Intensive Margins of the Hurdle Rate Buffer

Panel A: Extensive Margin of Buffer

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Linear Probability Model				Logit		
Cost of Capital	-0.027***	-0.028***	-0.029***	-0.029***	-0.025***	-0.026***	-0.028***
	(0.003)	(0.003)	(0.003)	(0.004)	(0.003)	(0.003)	(0.004)
Beta Volatility		0.031***	0.031***	0.013		0.031***	0.015
		(0.011)	(0.011)	(0.011)		(0.011)	(0.011)
Sales Volatility		0.010	0.009	0.015		0.010	0.015
		(0.015)	(0.015)	(0.016)		(0.016)	(0.016)
Large Firm		-0.077**				-0.085**	-0.067
		(0.036)				(0.039)	(0.044)
Public Firm		0.079*				0.065*	0.045
		(0.044)				(0.038)	(0.038)
Has Credit Rating		-0.012				-0.016	-0.061**
		(0.029)				(0.030)	(0.029)
Firm-Level Optimism				0.033*			0.033*
				(0.018)			(0.018)
US Economy Optimism				-0.001			-0.015
				(0.018)			(0.018)
Observations	1,232	1,232	1,232	947	1,232	1,232	947
R-squared	0.075	0.086	0.122	0.129	0.068	0.080	0.105
Survey Quarter FE			Yes	Yes			
Size FE			Yes	Yes			
Ownership FE			Yes	Yes			
Credit Rating FE			Yes	Yes			

Panel B: Intensive Margin of Buffer

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	All CFOs				Buffer > 0			
Cost of Capital	-0.307***	-0.311***	-0.314***	-0.319***	-0.194***	-0.181***	-0.180***	-0.166**
	(0.045)	(0.045)	(0.045)	(0.058)	(0.054)	(0.055)	(0.057)	(0.069)
Beta Volatility		0.432***	0.431***	0.325*		0.295*	0.289*	0.307
		(0.148)	(0.145)	(0.175)		(0.165)	(0.166)	(0.196)
Sales Volatility		0.395**	0.375**	0.561***		0.438***	0.420***	0.580***
		(0.163)	(0.167)	(0.173)		(0.151)	(0.155)	(0.171)
Large Firm		-1.246***				-1.033**		
		(0.391)				(0.448)		
Public Firm		-0.040				-0.610		
		(0.462)				(0.521)		
Has Credit Rating		-0.142				-0.077		
		(0.292)				(0.373)		
Firm-Level Optimism				0.520**				0.436*
				(0.210)				(0.235)
US Economy Optimism				-0.371				-0.476
				(0.267)				(0.313)
Observations	1,232	1,232	1,232	947	957	957	957	750
R-squared	0.058	0.083	0.099	0.120	0.021	0.048	0.064	0.083
Survey Quarter FE			Yes	Yes			Yes	Yes
Size FE			Yes	Yes			Yes	Yes
Ownership FE			Yes	Yes			Yes	Yes
Credit Rating FE			Yes	Yes			Yes	Yes

This table explores the determinants of the extensive and intensive margins of the hurdle rate buffer. In Panel A, we focus on the extensive margin (i.e., the dependent variable is the binary variable Uses Buffer {0,1}). The first four columns display linear probability models (OLS) and columns 5-7 display logistic regressions. In Panel B, we focus on the intensive margin. The variables Beta Volatility, Sales Volatility, Firm-Level Optimism and US Economy Optimism are standardized to mean zero, unit variance. All variables are defined in detail in Table B.1. In columns 4 and 7 of Panel A and columns 4 and 8 of Panel B, the number of observations drops because we do not have optimism forecasts from every respondent. Standard errors are clustered at survey industry × survey quarter and displayed in parentheses below the coefficient. The R^2 in columns 5-7 of Panel A is the pseudo-R^2 from the logistic regression. ***, ** , * denote significance at 1%, 5%, 10%.

Table 3.4: The Buffer, the Cost of Capital and Beta for Publicly Listed Companies.

	(1)	(2)	(3)	(4)	(5)	(6)
	Extensive Margin			Intensive Margin		
Cost of Capital	-0.032***	-0.032***	-0.034***	-0.432**	-0.439**	-0.532***
	(0.010)	(0.010)	(0.009)	(0.202)	(0.197)	(0.191)
Firm Beta	0.078***	0.077***	0.067**	1.001***	0.967***	1.252***
	(0.027)	(0.027)	(0.027)	(0.342)	(0.331)	(0.330)
Beta Volatility		0.012	0.005		0.422	0.334
		(0.020)	(0.019)		(0.278)	(0.274)
Sales Volatility		0.030	0.037		0.599**	0.499*
		(0.030)	(0.032)		(0.256)	(0.286)
Observations	186	186	186	186	186	186
R-squared	0.057	0.064	0.133	0.085	0.113	0.188
Survey Quarter FE			Yes			Yes
Size FE			Yes			Yes
Ownership FE			Yes			Yes
Credit Rating FE			Yes			Yes

This table explores how the firm's cost of capital affects the use and size of the buffer for publicly traded firms that appear in the CFO survey. In each specification, we control for the firm's estimated CAPM beta, in order to hold constant how the firm's underlying cash flow risk simultaneously affects the buffer and the cost of capital (via its effect on beta). Firm Beta is the firm's estimated CAPM beta; Beta Volatility and Sales Volatility are the industry level measures as shown in Table B.1; all three variables are standardized to mean zero, unit variance. Standard errors are clustered at survey industry × survey quarter and displayed in parentheses below the coefficient. ***, ** , * denote significance at 1%, 5%, 10%.

Table 3.5: The Buffer and Customer Concentration

	(1)	(2)	(3)	(4)	(5)	(6)
	Extensive Margin			Intensive Margin		
Customer Concentration	0.008	0.004	0.004	0.416***	0.295***	0.296**
	(0.013)	(0.010)	(0.012)	(0.140)	(0.082)	(0.107)
Cost of Capital	-0.027***	-0.027***	-0.029***	-0.313***	-0.320***	-0.317***
	(0.003)	(0.003)	(0.003)	(0.045)	(0.049)	(0.043)
Sales Volatility		0.007	0.008		0.369**	0.315*
		(0.010)	(0.010)		(0.133)	(0.143)
Beta Volatility		0.030**	0.030**		0.435**	0.399**
		(0.011)	(0.012)		(0.154)	(0.163)
Observations	1,232	1,232	1,232	1,232	1,232	1,232
R-squared	0.075	0.080	0.122	0.064	0.075	0.102
Survey Quarter FE			Yes			Yes
Size FE			Yes			Yes
Ownership FE			Yes			Yes
Credit Rating FE			Yes			Yes

This table explores how the bargaining power of a firm's customers affects the buffer. If a company's customers are concentrated, all else equal the company has less bargaining power; and vice versa. Our measure of customer bargaining power is customer concentration: the HHI of a firm's sales to corporate customers (Patatoukas, 2012). We aggregate to NAICS-3 by year level and match to our CFO survey data by their industries. Columns 1-3 focus on the extensive margin; columns 4-6 focus on the intensive margin. The variables Customer Concentration, Beta Volatility and Sales Volatility are standardized to mean zero, unit variance. All variables are defined in detail in Table B.1. Standard errors are clustered at survey industry × survey quarter and displayed in parentheses below the coefficient. ***, ** , * denote significance at 1%, 5%, 10%.

Table 3.6: The Buffer and Price Markups in B2B Industries

	(1)	(2)	(3)	(4)	(5)	(6)
	Extensive Margin			Intensive Margin		
Markup	-0.031***	-0.024**	-0.026*	-0.698***	-0.527***	-0.559***
	(0.011)	(0.011)	(0.013)	(0.148)	(0.162)	(0.163)
Cost of Capital	-0.027***	-0.028***	-0.028***	-0.326***	-0.328***	-0.324***
	(0.003)	(0.003)	(0.003)	(0.047)	(0.047)	(0.047)
Sales Volatility		0.003	0.004		0.284*	0.222
		(0.015)	(0.015)		(0.145)	(0.153)
Beta Volatility		0.027**	0.028**		0.426**	0.408**
		(0.013)	(0.013)		(0.167)	(0.165)
Observations	1,066	1,066	1,066	1,066	1,066	1,066
R-squared	0.083	0.087	0.120	0.079	0.086	0.113
Survey Quarter FE			Yes			Yes
Size FE			Yes			Yes
Ownership FE			Yes			Yes
Credit Rating FE			Yes			Yes

This table explores how a firm's level bargaining power affects the buffer. The proxy for a firm's bargaining power in this table is the price markup of firms operating in non-consumer-facing industries, i.e., firms whose operations are predominantly (business-to-business) B2B. Higher price markups on a firm's customers would be a direct consequence of bargaining power, all else equal. Our measure of price markup is the "accounting" markup from Baqaee and Farhi (2020b), which we aggregate to NAICS-3 by year level and match to our CFO survey data by their industries. Because our model makes no prediction on the relation between the buffer and markup for consumer-facing firms, we remove consumer-facing industries from the analysis. Specifically, we follow Gofman et al. (2020) and define an industry as consumer-facing if it falls in GICS sector "Consumer Discretionary" or "Consumer Staples" (GICS codes 25 and 30, respectively); 166 of our 1232 observations are firms operating in consumer-facing industries. In Appendix Table B.5 Panel A, we analyze the relation between the buffer and price markups for the full sample. Columns 1-3 focus on the extensive margin; columns 4-6 focus on the intensive margin. The variables Markup, Beta Volatility and Sales Volatility are standardized to mean zero, unit variance. All variables are defined in detail in Table B.1. Standard errors are clustered at the survey industry × survey quarter and displayed in parentheses below the coefficient. ***, ** , * denote significance at 1%, 5%, 10%.

Table 3.7: Differences in the Buffer: By Broad Industry Category and Asset Tangibility

Panel A: Broad Industry Categories

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Extensive Margin				Intensive Margin			
Manufacturing, Construction, Mining, Energy	0.073***				0.793***			
	(0.023)				(0.271)			
Finance		-0.058*				-1.149**		
		(0.034)				(0.454)		
Services, Comms, Media			-0.031				-0.025	
			(0.037)				(0.568)	
Healthcare, Pharma, Tech				-0.027				0.117
				(0.036)				(0.621)
Cost of Capital	-0.029***	-0.029***	-0.028***	-0.028***	-0.316***	-0.322***	-0.314***	-0.315***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.044)	(0.045)	(0.045)	(0.046)
Beta Volatility	0.030***	0.029**	0.031***	0.031***	0.421***	0.399***	0.432***	0.431***
	(0.011)	(0.011)	(0.011)	(0.011)	(0.145)	(0.143)	(0.145)	(0.145)
Sales Volatility	0.012	0.005	0.010	0.010	0.404**	0.301*	0.376**	0.370**
	(0.014)	(0.014)	(0.015)	(0.015)	(0.163)	(0.162)	(0.163)	(0.171)
Observations	1,232	1,232	1,232	1,232	1,232	1,232	1,232	1,232
R-squared	0.128	0.123	0.122	0.122	0.104	0.103	0.099	0.099
Survey Quarter FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Size FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ownership FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Credit Rating FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Panel B: Asset Tangibility

	(1)	(2)	(3)	(4)	(5)	(6)
	Extensive Margin			Intensive Margin		
Tangibility	0.032**	0.035**	0.037***	0.287*	0.389***	0.460***
	(0.014)	(0.014)	(0.014)	(0.145)	(0.145)	(0.148)
Cost of Capital	-0.027***	-0.027***	-0.029***	-0.308***	-0.294***	-0.315***
	(0.003)	(0.003)	(0.003)	(0.044)	(0.042)	(0.044)
Sales Volatility			0.015			0.450***
			(0.015)			(0.165)
Beta Volatility			0.029**			0.406***
			(0.011)			(0.140)
Observations	1,232	1,232	1,232	1,232	1,232	1,232
R-squared	0.080	0.088	0.129	0.061	0.085	0.106
Size FE		Yes	Yes		Yes	Yes
Survey Quarter FE			Yes			Yes
Ownership FE			Yes			Yes
Credit Rating FE			Yes			Yes

This table explores how the buffer varies across broad industry categories and across industry tangibility. In Panel A, we regress the buffer on indicator variables that cover the different industries in our sample. Columns 1-4 focus on the extensive margin; columns 5-8 focus on the intensive margin. Panel B explores how the buffer varies by industry tangibility (i.e., the ratio of property, plant and equipment to total fixed assets). Columns 1-3 focus on the extensive margin; columns 4-6 focus on the intensive margin. The variables Asset Tangibility, Beta Volatility and Sales Volatility are standardized to mean zero, unit variance. All variables are defined in detail in Table B.1. Standard errors are clustered at survey industry × survey quarter and displayed in parentheses below the coefficient. ***, ** , * denote significance at 1%, 5%, 10%.

4

Corporate Flexibility in a Time of Crisis

4.1 Introduction

Firms constantly respond to challenging circumstances, often triggered by unexpected shocks. A large body of research shows how financial flexibility allows firms to adapt to adversity. For example, Chodorow-Reich (2014) shows that less financially constrained businesses maintained stronger employment during the 2008-9 Global Financial Crisis. More recently, Acharya and Steffen (2020) show that companies raised cash to prevent distress following the outbreak of COVID-19. While the importance of financial flexibility has been well documented, the unprecedented challenges brought on by the 2020 health crisis highlight the importance of additional dimensions of corporate flexibility, such as the ability to perform work remotely (e.g., Barrero et al., 2021; Papanikolaou and Schmidt, 2022), in helping firms manage difficult circumstances.

We conduct a series of CFO surveys to gather companies' internal plans, which we use to study the role of corporate flexibility — the ability of firms to adjust and adapt — in response to the COVID-19 crisis. While existing papers investigate a number of important aspects of the impact of COVID-19, our contribution is unique in providing a comprehensive investigation of firms' real-side decisions, including employment and capital expenditures. Our series of surveys allow us to analyze companies' planning in real time as the COVID crisis hit as well as their

long-term outlook for the post-COVID recovery. We confirm the key findings with archival data. We examine *multiple dimensions* of flexibility and highlight the importance of three margins: 1) financial flexibility, which captures the well-studied observation that financial resources are important for supporting adjustments in firms' activities; 2) workplace flexibility, which refers to firms' abilities to accommodate remote work in order to maintain social distancing; and 3) investment flexibility, which reflects whether firms can modify the timing of their capital spending in response to changing conditions. In this paper, we show i) how each of these dimensions plays a role in shaping corporate planning in the COVID-19 crisis and ii) how they interact.

Our primary data come from a survey conducted from mid-February to mid-April 2020, capturing in real time the responses of US companies to the sudden COVID-19 outbreak; in particular, their plans for hiring and investing. Companies in our sample represent all sectors and 47 of the 50 states. They include large, medium, and small firms, as well as public and private entities. We continued to survey CFOs in June, September, and December of 2020. These additional survey waves confirm the trajectory of business plans for 2020; and importantly, provide information on companies' long-term planning for the post-COVID world.

The strategy we use to study various dimensions of corporate flexibility is straightforward. Our financial flexibility measure captures CFOs' (survey-based) assessments, and reflects both the availability of internal funds and access to external financing. For workplace flexibility, we identify the extent to which employees are able to work remotely (cf. Papanikolaou and Schmidt, 2022; Dingel and Neiman, 2020). For investment flexibility, we use survey-based information on firms' abilities to adjust the timing of their capital expenditures; that is, whether to delay or accelerate spending. Using O*NET and BEA data, we also create a proxy to control for businesses' exposure to reduced demand based on social interactions consumers face in the purchase process (i.e., both direct exposure and indirect exposure through downstream industries).

To organize our analysis, we start with a simple model framework that illustrates how different aspects of corporate flexibility affect firms' real decisions. More financial flexibility, such as having more cash available, relaxes firms' funding constraints and supports more employment and investment. Workplace flexibility helps firms stay productive during the COVID emergency,

reducing health risks from traditional in-person work and boosting employment. Investment flexibility allows firms that are experiencing adverse conditions and production difficulties to defer capital expenditures during COVID (or accelerate investment if conditions are favorable). As we detail below, our model focuses on how multiple dimensions of corporate planning are affected following the onset of the pandemic, complementing recent work by Acharya et al. (2021) on firms' precautionary tradeoff between maintaining financial efficiency and operational resiliency.

Our first set of empirical tests looks at firms' real decisions as the pandemic hit, conditional on their financial, workplace, and investment flexibility. We document key determinants of companies' planning for employment and capital spending during the COVID outbreak. To begin, we find that firms with more financial flexibility planned for higher employment and capital expenditure growth in 2020, a result that resonates with the positive impact of financial flexibility on stock returns reported by existing studies (e.g., Fahlenbrach et al., 2021; Ramelli and Wagner, 2020). Notably, we find that the impact of financial flexibility on employment is significantly stronger for firms with more fixed costs. This evidence aligns with economic intuition: if a company's costs are entirely fixed, then revenues decline in a crisis but costs do not. In this case, having access to cash can be especially important for covering fixed costs. In contrast, if a company's costs are entirely variable, then costs will also decrease as production activities and revenues fall in a crisis, alleviating the need for extra cash.

In addition to studying financial flexibility, a novel aspect of our paper is simultaneously exploring the roles played by two additional forms of flexibility, as well as interactions between all three forms of flexibility, as companies managed the COVID crisis. We find that higher workplace flexibility is associated with significantly higher planned employment growth. This positive effect of workplace flexibility echoes findings in the literature (see, e.g., Papanikolaou and Schmidt, 2022; Favilukis et al., 2020). At the same time, higher workplace flexibility does not appear to directly boost capital spending plans, which as we discuss later, suggests that remote work is likely to make traditional capital investment less relevant. We also explore the role of investment flexibility and its interaction with workplace flexibility. We find that workplace flexibility affects how firms use investment flexibility: companies with a flexible workplace can operate relatively

smoothly during the crisis and exploit higher investment flexibility to *increase* capital spending. In contrast, companies with low workplace flexibility experience unfavorable conditions and use higher investment flexibility to *reduce* — or possibly postpone — capital spending. We confirm that these results hold over time based on realized outcomes from subsequent surveys. We also provide external validation for our survey results using realized Compustat data. Finally, we consider additional types of flexibility related to workplace arrangements (e.g., part-time workers, scheduling autonomy, unionization) and adjustment costs (Gu et al., 2019, 2021), but do not find significant effects associated with these margins in the COVID crisis.

We perform further analyses to characterize the extent to which the above results reflect the unique challenges of the 2020 pandemic. Among other tests, we compare our findings during the COVID-19 crisis to the economic forces at play during the Global Financial Crisis using CFO survey data from Campello et al. (2010). We first show that financial flexibility appears to exert similar impact on employment and investment plans in both crises. We then turn to the analysis of workplace flexibility, noting that the physical environment and logistics of the workplace have evolved significantly in recent years.[1] Our results show that workplace flexibility played *no role* in firms' decisions during the 2008 financial crisis, while it is *central* in the 2020 health crisis. Likewise, our tests do not indicate that firms significantly exploited their investment flexibility during the 2008 crisis — at least not in tandem with their workplace flexibility. As external validation for these results, we use historical Compustat and BLS data to show that workplace flexibility did not significantly affect firm employment or investment in the 15 years before 2020. Our extensive checks confirm that our central findings are new to the COVID-19 crisis.

A final set of tests investigate companies' planning for the post-COVID world. We find that firms' responses to the pandemic may have accelerated changes in the *nature* of investment. Notably, firms with high workplace flexibility expect employment to recover more quickly and capital spending more slowly; these firms also expect remote work to persist for longer. Our evidence

[1] Barrero et al. (2021) point out the proportion of employees who primarily worked from home had grown from 0.8% in 1980 to 2.4% in 2010, reaching 4% in 2018. Around 40% of working age individuals were working from home in May 2020 (see also Bick et al., 2021; Brynjolfsson et al., 2020). See Bloom et al. (2014) and Mas and Pallais (2017) for pre-COVID studies on remote work.

suggests that this segment of the economy may steer away from investing in traditional physical assets, instead investing in the workforce and intangible assets that facilitate flexible collaboration. We also find that large firms and firms with lower workplace flexibility are inclined to implement automation to reduce their dependence on labor. This link between workplace flexibility and automation is new and can contribute to a persistent impact of workplace flexibility on employment decisions. Accordingly, we also investigate which types of workers are most affected by automation and find that low-skill workers are more likely to be displaced. In particular, low workplace flexibility firms show a stronger tendency to displace low-skill workers. While these steps may help firms better withstand future health crises or other workplace disruptions, the prospect of increased automation that we document may pose new, long-term challenges for many workers in the economy.

4.1.1 Literature Review

Our paper builds on several strands of research. First, an influential literature analyzes the impact of firms' financial capacity, such as the effects of financial constraints (e.g., Fazzari et al., 1988; Whited and Wu, 2006). Recent work has emphasized the role of financial flexibility in the COVID-19 crisis (as many firms' cash flows suddenly dried up), including how firms build cash holdings to boost financial flexibility (Acharya and Steffen, 2020) and how financially flexible firms experience better stock returns (Fahlenbrach et al., 2021; Ramelli and Wagner, 2020; Ding et al., 2021). We add to this literature by documenting the role of financial flexibility on firms' *real-side* planning for employment and investment. We further show that the impact of financial flexibility is contingent on the firm's cost structure: the effects of financial constraints are particularly strong among firms with high fixed costs.

Second, our work relates to recent research on work from home. Several papers document the work-from-home waves that occurred in 2020 and suggest that work from home will persist in the future (Bick et al., 2021; Barrero et al., 2021). Favilukis et al. (2020) and Papanikolaou and Schmidt (2022) find stronger stock performance among firms in industries with higher work-from-home capability, and Pagano et al. (2021) find higher stock returns for industries with better social dis-

tancing in general. Our contribution in this dimension is to utilize detailed firm-level data on real decisions about employment and investment, analyze the interactive role of workplace flexibility in connection with other forms of flexibility, and shed light on firms' decisions both during the COVID-19 outbreak and over the long-term. Our results show that workplace flexibility was not only a first-order determinant of employment decisions during the health crisis, but it may play a key role for the long-term transformation of both the nature of investment (firms with high workplace flexibility shift away from traditional capital expenditures) and the nature of the workforce (firms with low workplace flexibility adopt automation to replace labor).

Several existing papers survey firms about the effects of COVID-19. The majority of those surveys focus on small businesses (Bartik et al., 2020; Alekseev et al., 2020; Bloom et al., 2021); Barrero et al. (2020) cover a range of firm sizes, but focus on the reallocation angle of the COVID shock. Focusing on public companies, Hassan et al. (2021) perform textual analyses of conference calls and document that executives discuss disruptions of production and operations more during the COVID-19 pandemic than in previous disease outbreaks. Our data are inclusive of larger corporations that play a measurable role in their industries and the larger economy. We use our survey data to perform a comprehensive analysis of firms' *ex ante* real decisions during the COVID crisis, shedding light on key outcomes regarding hiring, capital investment, workplace arrangements, and automation. We also confirm our key findings using *ex post* realizations of firms' employment and capital spending based on Compustat data. Finally, there is an important literature analyzing firms' stock returns at the onset of the COVID-19 crisis, including several studies discussed above. We provide a detailed comparison of our results relative to those in the stock returns literature in Section 4.4.3.[2]

Taken together, our work adds to the literature for the new insights it brings into how corporate flexibility facilitates the real decisions firms make in responding and adapting to crises. In particular, our analysis advances knowledge into how firms jointly operate across *multiple margins* (financial, workplace, and investment flexibility), as well as their interactions, during

[2] Another set of studies focus on the impact of government policies. For instance, Chetty et al. (2020), Granja et al. (2020), and Balyuk et al. (2021), among others, examine the effects of the Paycheck Protection Program (PPP).

the COVID crisis. Our study complements recent work by Acharya et al. (2021) on the interface between firms' financial and operational margins.

4.2 Data and Summary Statistics

We describe our data sources and present summary statistics in this section. We describe the main variables covered by our CFO surveys in Section 4.2.1 and variables from other sources in Section 4.2.2. Internet Appendix C.3provides information on the construction of each variable used in the paper.

4.2.1 CFO Survey Data

Our baseline data source is the Global Business Outlook survey of US CFOs conducted by Duke University in the first quarter of 2020. The survey questionnaire is available at https://cfosurvey.fuqua.duke.edu/release/. This survey provides timely information about how firms respond to the sudden arrival of the COVID-19 crisis. We sent out e-mail invitations for this survey starting on February 11, 2020, before the escalation of the spread of the novel coronavirus across the US. This survey round closed on April 10, 2020. Because the timing is centered on March, we refer to it as the "March 2020" survey. We obtained survey responses from 520 CFOs. The overall response rate is 19.5%, which is high compared with typical surveys of executives and investors.[3]

Figure 4.1 summarizes key characteristics of the respondent firms and shows that the sample includes a wide variety of company types. Panel A shows that sample firms are spread across several industries, including services and manufacturing. Panel B shows that the sample includes large firms (revenue over $1 billion), as well as "middle market" (revenue between $100 million and $1 billion) and small firms. Panel C shows that about half of our responses were received before mid-March, when there were still few reported COVID-19 cases in the US. The other half of the responses were received after mid-March, following the national COVID-19 emergency declaration. Figure C.1 shows that the composition of firms is also similar among responses in

[3] The response rate for prior Duke CFO surveys was about 9% (Graham and Harvey, 2001b; Ben-David et al., 2013). Gompers et al. (2020) survey private equity investors and obtain a response rate of 23%. Giglio et al. (2021) survey Vanguard investors and obtain a response rate around 4%.

different survey weeks.

To provide more context for our analysis, we also compare our sample to Compustat firms based on 2019 financial information. Figure 4.2 shows the firm size distributions in both samples. Specifically, below each panel title, we display the fraction of our sample and Compustat that are within each revenue bracket. Our sample contains proportionally more small firms than Compustat, given that we include many private firms, but our sample also includes a reasonable number of medium-sized and large companies. Each panel also shows quartiles of employment within each revenue bracket. The comparison shows that, within each revenue bin, our sample largely captures the distribution of employment of Compustat firms. Taken together, our sample appears well-suited for studying the broad cross-section of companies.

The March 2020 survey asked CFOs about their projected growth in revenue, employment (domestic full-time employees), and capital expenditures (spending on structures and equipment) in 2020. In addition, to measure financial flexibility, we asked CFOs to assess the level of financial flexibility their firms have: "About how much financial flexibility would you say your company has right now? (0-None, 1-A little, 2-3-4-Moderate, 5-A lot)." We classify a firm as having financial flexibility if they answered 2 or greater. As we verify in Table C.1, this measure of financial flexibility captures both the abundance of internal funds and the ability to access external financing.

To measure investment flexibility, we use information for the 636 US companies that responded to the Duke Global Business Outlook survey conducted in March 2019 (before the COVID-19 crisis). It is not an easy task to gauge how flexible a firm's investment spending process is, but our survey instrument provides important insight into this issue. In particular, the March 2019 survey collected data on firms' flexibility in investment implementation by asking, "How flexible is the speed at which you complete your largest capital investment project? (0-Very flexible; 1-Flexible; 2-Somewhat flexible; 3-Neutral; 4-Somewhat inflexible; 5-Inflexible; 6-Very inflexible)." We classify a March 2019 firm as having high investment flexibility if the response is 0 or 1. We construct an industry-level measure of investment flexibility by calculating the percentage of firms with high investment flexibility at the four-digit NAICS level. This allows us to apply the 2019 measure of investment flexibility to the entire 2020 sample. We verify that this attribute has

an important industry component: the R^2 from four-digit NAICS fixed effects is 0.45. In subsequent analyses, we also show that our industry-level measure of investment flexibility produces consistent results in different samples (e.g., March 2020 survey, Compustat firms), which also suggests that it captures intrinsic characteristics of different industries. Conceptually, our investment flexibility measure captures flexibility in the *timing* of investment, which is especially relevant for firms' responses to a sudden crisis. This measure is novel compared to prior work on investment adjustment costs, which have focused on costs that depend on the magnitude of investment (as summarized by Cooper and Haltiwanger, 2006). For illustrative purposes, Table C.2 displays the measure at the two-digit NAICS level. While there is important variation at finer industry levels, the two-digit measure provides context for interpreting our results.

Following the March 2020 survey, we conducted additional surveys in June, September, and December 2020. We did so in collaboration with the Federal Reserve Banks of Atlanta and Richmond.[4] The September survey asked firms when they expect various labor and spending outcomes to return to pre-COVID levels, and the December survey explored automation.

4.2.2 Other Data

We collect data from several other sources to enhance our analysis. The external datasets measure firm attributes at the industry level, and we match them with firms in our CFO surveys based on their industries.[5]

For workplace flexibility, we collect data on employees' ability to work remotely by calculating the fraction of employees in each industry who can and do work from home using the American Time Use Survey (ATUS) (Papanikolaou and Schmidt, 2022; Hensvik et al., 2020; Alon et al., 2020). This measure is available for each four-digit NAICS code. We also perform additional tests using the fraction of employees in each industry who can work from home constructed by

[4] Collectively, nearly 650 firms responded to the June, September, and December surveys, with responses relatively evenly split across the three quarters. These surveys are publicly available at https://www.richmondfed.org/research/national_economy/cfo_survey/data_and_results. Our analysis of these data and conclusions do not necessarily reflect the views of the Federal Reserve Bank.

[5] We know the industry codes of public firms directly. For private firms, we use the company name provided by the CFO to infer the firm's industry using historical business data from services such as Infogroup, and from survey-based responses that list the names of other firms in their industries.

Dingel and Neiman (2020), which uses O*NET data and is available for two- and three-digit NAICS codes.[6] Traditionally, workplace flexibility has been examined in labor economics from the employee's perspective (e.g., remote work is a potential way to balance work and family). Since the COVID-19 outbreak, workplace flexibility also became central from the firm's perspective as remote work helps companies adapt to health risks and social distancing.

The workplace flexibility measures from ATUS and O*NET are both constructed using data prior to the COVID-19 crisis, and we cross check these *ex ante* measures with the *ex post* prevalence of remote work reported by the BLS every month since May 2020. For each industry corresponding to roughly two-digit NAICS codes, the BLS data show the fraction of employees who worked remotely in the last four weeks due to COVID-19. We find the *ex ante* measures are around 80% correlated with the BLS measure, which confirms their informativeness. Figure C.7in the Internet Appendix presents a scatter plot demonstrating the high correlation between the pre-COVID ATUS measure (x-axis) and the July 2020 BLS measure (y-axis). While the overall prevalence of work from home increased substantially after the COVID outbreak, the cross-sectional relation across industries stayed constant.

We also collect data to control for the impact of the COVID-19 crisis on *customer demand*. In a pandemic, customers reduce activities that require social contact. To control for this factor, we follow the literature (e.g., Koren and Pető, 2020; Pagano et al., 2021) and measure the degree of "in person" transactions in an industry by looking at how much employees interact with customers to provide goods and services. We use the O*NET Work Activities survey, which asks workers how important "performing for people or dealing directly with the public" is to their job. We aggregate this to the industry level (four-digit NAICS) and term it "Direct Customer Interactions."

We note that the above measure only appropriately proxies the effects of the COVID-19 demand reduction for firms that are directly consumer-facing. It ignores the fact that this reduction in customer demand may travel along the supply chain. For instance, although Boeing

[6] The ATUS and O*NET (Dingel and Neiman, 2020) work-from-home measures use similar sources but their construction procedures contain important differences. The ATUS measure captures whether workers can work from home (and have done so in the past). The O*NET measure captures the nature of the work that employees perform at the occupation level. Internet Appendix C.3discusses these measures in detail.

does not score high on direct customer interactions, Boeing's demand was negatively affected by the health crisis because its main downstream industry (airlines) has high direct customer interactions. Hence, we also construct an "Indirect Customer Interactions" measure by taking the weighted average of "Direct Customer Interactions" among a given firm's *downstream industries*, where the weights are computed using the fraction of sales to different downstream industries according to the Input-Output Table of the BEA. Industries such as retail and travel rank high for direct customer interactions and low for indirect customer interactions, whereas manufacturing industries that supply to them rank high for indirect customer interaction and low for direct customer interactions. Fahlenbrach et al. (2021) analyze how customer interactions filter through the supply chain and affect stock returns, and our indirect measure is similar in construction. The "Customer Interactions" variable we use in our empirical analysis is a weighted average of the direct and indirect measures.[7]

4.2.3 Summary Statistics

Table 4.1 reports summary statistics for the March 2020 survey and Table C.5in the Internet Appendix does the same for the three subsequent surveys from June, September and December of 2020.[8]

Measures of Flexibility. For financial flexibility, about 20% of firms are classified as having low financial flexibility. Concerning workplace flexibility, for the average firm, about 25% of employees in its industry can work from home (and have done so in the past) according to the ATUS data, which we use as our primary measure (four-digit NAICS code level). At the same time, 45% of employees can (in principle) work from home based on the data of Dingel and Neiman (2020), which we use as an additional measure (two-digit NAICS code level). For investment flexibility, on average, about 26% of firms in an industry indicate that they can adjust the speed of capital investment flexibly.

[7] In particular, Customer Interactions$_i$ = C_i(Direct$_i$)+(1−C_i)(Indirect$_i$), where C_i is the industry's share of output that belongs to "Personal Consumption Expenditures" in the BEA's Input-Ouput Table (i.e., directly selling to consumers) and 1 − C_i is the share of sales that belong to intermediate goods.

[8] Our key industry-level measures are available at https://github.com/jwb4335/corporate_flexibility.

Figure C.6in the Internet Appendix shows pairwise variable correlations from the March 2020 survey. Workplace flexibility and financial flexibility are not highly correlated; workplace flexibility and investment flexibility are weakly positively correlated.

CFO Outlook. Table 4.1 shows that, in the March 2020 survey, CFOs expect 4.6% annual revenue growth. They also expect a modest average level of employment growth and capital expenditure growth for the year. As we show using Figure C.2, CFOs' revenue expectations changed substantially from early March through early April, as the severity of the pandemic escalated. The projections were between 5% to 10% in early March, but collapsed to approximately 0% by late March and early April, staying steady in subsequent surveys (though respondents change from survey to survey).[9] Figure C.2 also shows a high degree of alignment between the CFO forecasts in our data and equity analyst forecasts reported in IBES (studied by several recent papers, e.g., Landier and Thesmar, 2020; Hong et al., 2021). This consistency indicates that firms in our survey are representative of firms in general. Overall, COVID-19 is a large shock and the revenue impact for 2020 is substantial, comparable in magnitude to the negative revenue impact firms experienced in 2009. Table C.5in the Internet Appendix shows that CFOs, on average, expect revenue growth through 2021 to return to close to 10%. Although CFOs expect a relatively speedy revenue recovery, as we show in Sections 4.4 and 4.5, this crisis affects the very organizational structure of firms and CFOs expect these long-term effects to persist after revenues recover.

Firm Heterogeneity. Since the onset of COVID-19, it has become apparent that the severity of the shock was different across firms. The March 2020 survey asked CFOs about their firms' exposure to risk from COVID-19. Table C.6in the Internet Appendix examines the determinants of this self-assessed risk exposure. We find that lower workplace flexibility and lower investment flexibility are associated with higher perceived COVID risk exposure. Lower financial flexibility is also associated with higher COVID risk exposure, but the statistical relation is weak, suggesting that CFOs do not necessarily consider the 2020 crisis to be financial in nature. Higher customer interactions (direct or indirect) are also associated with higher COVID risk in the view of CFOs.

[9] In the June 2020 survey, CFOs indicated that COVID-19 would directly have a 10% negative impact on their firms' 2020 revenue growth, which aligns with the expected 10% revenue growth reduction as of the March survey.

While subjective, this COVID-related risk assessment gives insight into CFOs' perceptions of the multifaceted challenges brought by COVID-19. In particular, the alleviating role of corporate flexibility provides a basis for our exploration into its effects on real decisions.

4.3 Conceptual Framework: The Role of Flexibility during the COVID-19 Outbreak

As discussed above, the challenges brought about by the COVID-19 crisis are multifaceted. Hence, we propose a basic conceptual framework to study how multiple dimensions of corporate flexibility affect firms' real decisions during a crisis. This simple framework helps us organize our subsequent empirical work. We discuss in turn its building blocks.

First, it is well understood that financing constraints can be relevant for firms' decisions in a variety of crises (even crises that do not originate from the financial sector), as firms rely on financial resources to support their operations and avoid financial distress. Indeed, these issues have been the focus of several papers on firms' financial policies or stock performance during COVID-19 (see Acharya and Steffen, 2020; Fahlenbrach et al., 2021; Ramelli and Wagner, 2020). We refer to this margin as "financial flexibility." Our empirical measure summarizes firms' ability to access *both* internal and external funding, as explained in Section 4.2.1 (see also Table C.1).

Second, as many corporate executives highlight, workplace flexibility — the ability for employees to work from home — is a key issue during the COVID-19 crisis. This theme is also reflected in recent academic studies that measure and analyze the prevalence of work from home (Dingel and Neiman, 2020; Barrero et al., 2021; Papanikolaou and Schmidt, 2022; Alon et al., 2020; Hensvik et al., 2020). Workplace flexibility became critical as the pandemic unfolded, as it allowed for better social distancing practices and helped employees balance caring for family members as needed. Firms whose employees cannot easily work from home may need to adopt additional health protocols (even limit production capacity) to control infection risk and maintain social distancing at work. Accordingly, low workplace flexibility — the inability to work from home — could negatively affect firms and their productivity in the pandemic. Our work integrates this margin together with other levers that are relevant to companies.

Third, faced with the sudden pandemic outbreak, firms needed to consider their ability to adjust the timing of their investment projects. Naturally, the manner in which companies utilize investment flexibility should depend on the circumstances they face. In particular, firms experiencing favorable conditions can utilize greater flexibility to front-load investment. In contrast, firms experiencing unfavorable conditions due to the pandemic can utilize greater investment flexibility to delay capital spending during difficult times. As a result, we expect investment flexibility to interact with factors that determine whether firms face favorable or unfavorable conditions. As we demonstrate below, workplace flexibility is an important margin in determining whether a firm faces favorable operating conditions in the COVID crisis, and the degree of workplace flexibility modulates how firms use their investment flexibility. This analysis reveals the role of investment flexibility as a margin of adjustment for addressing a crisis; the conditional impact of investment flexibility also helps us identify which factors are important for shaping whether firms experience favorable or unfavorable conditions.

A Simple Model. We integrate the above observations into a simple two-period model. We do so by following an influential body of work that highlights that the core shock induced by COVID-19 is a supply shock (Guerrieri et al., 2021; Eichenbaum et al., 2021; Fornaro and Wolf, 2020; Acemoglu et al., 2021; Baqaee and Farhi, 2020a, 2021; Favilukis et al., 2020), where the magnitude of this shock is modulated by the degree of workplace flexibility. Guerrieri et al. (2021), for example, model the COVID-induced supply shock as "a fraction $\phi > 0$ of agents become unable to produce," so production changes from $Y_t = N_t$ to $Y_t = (1 - \phi)N_t$.[10] Favilukis et al. (2020) provide a model that examines several possible shocks associated with COVID. Those authors show that to match the empirical evidence it is important to model COVID as a productivity shock that hits a firm harder if the firm is less able to implement remote work. Correspondingly, we also model COVID as a productivity shock, where its impact is modulated by the degree of workplace flexibility. To the extent a supply shock induced by COVID may also translate into an aggregate demand shock (Guerrieri et al., 2021), we further allow for a demand shifter in the firm's production function

[10] One can interpret the shock as either only $(1 - \phi)$ fraction of agents can work at a given point of time, or all agents can work at $(1 - \phi)$ fraction of the pre-COVID capacity. Either interpretation is sufficient for the purposes of our analysis.

similar to Favilukis et al. (2020).

Period 1 in our model corresponds to the COVID outbreak, when firms respond to a sudden and unanticipated health crisis. In this period, firms are hit by a supply shock that forces them to operate at a fraction ϕ of normal productivity A (e.g., workers cannot come to production facilities, so the supply of labor and the use of capital both decline). Our model also allows firms to experience a shift in demand, indexed by B (this maps into the customer interaction control variable we include in our empirical analyses below). In response to these shocks, the firm plans for employment, and adjusts its previous investment plans for this period (pre-determined value K_1) by deferring Δ units of investment to Period 2. Period 2 corresponds to production after the COVID crisis.

The firm optimizes by choosing its labor (governed by L_1 and L_2) and capital (governed by Δ and K_2). For simplicity, capital depreciates fully across periods and the discount rate is normalized to one. The wage rate of workers is denoted by w. The firm solves the following maximization problem:

$$\begin{aligned} \max_{L_1,\Delta,L_2,K_2} \quad & B[\phi A L_1^{\alpha} + \phi A(K_1 - \Delta)^{\beta}] + AL_2^{\alpha} + A(K_2 + \Delta)^{\beta} \\ & - (wL_1 + wL_2) - (K_1 + K_2 + \xi\Delta^2) \\ & \text{s.t.} \quad wL_1 + K_1 - \Delta \leq C. \end{aligned}$$

In this setup, financial flexibility during the COVID crisis is represented by a financial constraint that requires spending on labor and capital to be less than available cash C (this can include both internal cash and external financing such as funding through credit lines). As discussed earlier, workplace flexibility modulates the magnitude of the supply shock ϕ during the pandemic. A firm with higher workplace flexibility and a greater ability to accommodate remote work can be relatively more productive in the health crisis (as reflected in higher ϕ). Finally, investment flexibility is captured by firms' ability to modify original plans for investment in period 1 (K_1): the firm can change the amount of investment by Δ and make up for it later in period 2, which comes with a cost $\xi\Delta^2$. This feature of our model is in line with evidence that firms establish their investment schedules ahead of time, but can modify them at some cost in light of changing circumstances (Lamont, 2000; Charoenwong et al., 2021). Firms with higher investment flexibility

have a lower ξ. We provide the derivations of model solutions and derivations of comparative statics in Internet Appendix C.2.

This model framework yields several predictions informing our empirical analyses. First, when the financial constraint is binding, both labor (L_1) and investment ($K_1 - \Delta$) in period 1 will increase with C (more cash available to support expenditures in period 1). This prediction verifies the importance of financial flexibility. Second, higher workplace flexibility (and therefore higher ϕ) is associated with higher employment in the pandemic (L_1).[11] Third, if the firm decides to defer investment given adverse conditions during COVID (i.e., $\Delta > 0$), then more investment flexibility will lead to more deferral (i.e., Δ is higher when ξ is smaller). If the firm decides to front-load investment (i.e., $\Delta < 0$) because it experiences relatively favorable conditions, then more investment flexibility will lead to more front-loading (i.e., Δ is more negative when ξ is smaller). Finally, for the demand shifter B, higher demand would generally boost employment and investment. The demand shifter can also interact with investment flexibility: firms with stronger (weaker) demand face favorable (adverse) conditions, and will use greater investment flexibility to front-load (defer) investment.

We note that as the COVID-19 outbreak is an unanticipated, unprecedented emergency, the workplace flexibility and investment flexibility measures we focus on reflect the nature of business operations in different industries, which affect companies' decisions following the onset of the crisis. In addition to the influence of external industry attributes, individual companies may also develop operational strategies and precautions in advance of crises. Recent work by Acharya et al. (2021) analyzes this dimension and investigates how financial hedging and operational hedging interact.

Other Forms of Flexibility. Our main analyses focus on three forms of flexibility that are central in the COVID-19 crisis. We also explore several other forms of flexibility and provide additional tests in Section 4.4.1. First, firms with higher fixed costs cannot easily reduce their operating

[11] One exception is that higher workplace flexibility may not boost employment when financial constraints are binding and firms choose to front-load investment (i.e., $\Delta < 0$). In this case, due to the binding financial constraints, labor competes with capital for financial resources. When ϕ is high, if sufficiently more resources are given to investment, employment could decline.

costs when negative shocks hit. Although fixed costs do not directly affect marginal returns from employment and investment, they can make financial resources (e.g., cash) especially valuable. In particular, when revenues fall during the COVID-19 outbreak, high fixed costs can drain firms' cash holdings and make the financial constraint more binding. Second, firms may differ in other dimensions of the workplace arrangement, such as the use of part-time *vs.* full-time workers and the ability to schedule work hours flexibly. Third, there can be other forms of adjustment costs of labor or capital, such as unionization and traditional capital adjustment costs. Our empirical work will consider these as well as other forms of flexibility.

4.4 Corporate Plans to Hire and Invest during the COVID-19 Crisis

Employment and capital expenditures are among the most important elements of corporate decision-making. Through our survey instrument, we are uniquely able to study CFO's *forward-looking plans* to hire and invest in *real time* amid the pandemic. We do so in this section. Towards the end of the section, we compare our COVID-19 findings to actions taken by firms during the 2008 Financial Crisis. We also validate our survey findings with subsequent surveys and realized archival data.

4.4.1 Corporate Flexibility in the 2020 Health Crisis

Basic Results on the Impact of Corporate Flexibility

In Table 4.2, we follow the framework introduced in the previous section and test the effects of financial flexibility, workplace flexibility, and investment flexibility on firms' employment and capital spending plans in real time as the COVID-19 crisis hit the US. We do so using CFOs' projections of employment and capital spending growth in 2020 from the March 2020 survey. Panel A presents the results from our main tests using the ATUS work-from-home measure at the four-digit NAICS level; Panel B does the same using the Dingel and Neiman (2020) work-from-home measure.

Results in Table 4.2 show that higher financial flexibility is associated with higher planned employment and capital expenditure growth in 2020. This is consistent with predictions discussed

above and prior findings on the impact of financial flexibility on corporate plans. All else equal, firms with low financial flexibility expect 7–9 percentage point lower growth of employment and capital expenditures in 2020. Notably, higher workplace flexibility is also associated with significantly higher projections of employment growth during the pandemic, consistent with our predictions. This result holds for both measures of workplace flexibility. Firms at the 75th percentile of the fraction of employees who can work from home expect 3–4 percentage point higher employment growth than those at the 25th percentile.[12] In untabulated tests, we find that this effect remains significant among the subset of firms with high financial flexibility, in line with predictions. Interestingly, higher workplace flexibility does not directly translate into higher projections of capital expenditure growth. As we discuss later, this evidence suggests that firms at which employees can work from home may be shifting towards new forms of investment and away from traditional capital expenditures.

Table 4.2 also shows that investment flexibility does not have a clear, unconditional impact on real decisions, in line with our predictions. As discussed above, how firms use investment flexibility should depend on whether they face favorable or unfavorable conditions, and we document this interactive result in the data shortly.

Overall, our results are robust to including controls of customer demand and firm size (log employees at the end of 2019), though neither of these variables is significant.[13] They are also robust to the inclusion of time (calendar week) fixed effects and state fixed effects (which absorb state-level variations in pandemic policies such as stay-at-home orders). Finally, they are robust to the inclusion of two-digit NAICS fixed effects, indicating that there is meaningful variation at the finer industry level.

[12] As shown in Table 4.1, the interquartile range of workplace flexibility is 0.3 for the ATUS measure and 0.5 for the Dingel and Neiman (2020) measure. The regression coefficients in Table 4.2 are between 0.08 and 0.1 for both measures. The difference between firms in the top and bottom quartile of workplace flexibility is between 0.3×0.1 = 0.03 and 0.5×0.08 = 0.04.

[13] Table 4.2 is based on the March 2020 survey data. In untabulated analysis, when we use all four quarters of our survey data, the customer interaction variable becomes significantly negatively related to employment, as expected. Thus, this negative relation between customer interaction and employment is driven by data from the second half of 2020. This conclusion is consistent with analysis of full-year Compustat data in Table 4.7, where again there is a negative association between customer interaction and employment. Detailed analysis available upon request.

Figure 4.3 depicts two of the insights from Table 4.2 via binscatter plots, displaying planned employment and capital expenditure growth plotted against workplace flexibility. In a related paper, Papanikolaou and Schmidt (2022) use industry-level monthly employment data from the BLS and report that total employment growth was higher from March to April in industries with more workplace flexibility. The evidence in Table 4.2 and Figure 4.3, Panel A, shows this important pattern in *ex ante* firm-by-firm plans gathered directly from CFOs, as they reveal firms' real-time decisions for full-year 2020. As shown in the table and figure, our data also allow us to study firms' investment decisions jointly with their employment decisions, which is important for a broader understanding of the effects of various levers of corporate flexibility.

The Conditional Impact of Investment Flexibility

We expand our analysis of corporate flexibility by unpacking the conditional nature of investment flexibility. As discussed in Section 4.3, we expect firms experiencing favorable *vs.* unfavorable conditions to use their investment flexibility differently. In terms of what determines favorable conditions, we find that workplace flexibility plays a key role, which we first illustrate in Figure 4.4. Panel A of Figure 4.4 shows that among firms with low workplace flexibility (which therefore face challenging operational conditions), those with high investment flexibility expect capital expenditures to fall by approximately 10% on average (indicating reductions or deferrals for firms that have the investment flexibility to do so), while those with low investment flexibility expect nearly 4% capital expenditure growth in 2020. In contrast, Panel B shows that among firms with high workplace flexibility (which therefore face more favorable operating conditions), those with higher investment flexibility plan to invest more during the pandemic. These patterns demonstrate that investment flexibility shapes firms' abilities to reduce *vs.* accelerate capital expenditures, and that this effect is conditional on workplace flexibility in an economically sensible way.

Table 4.3 characterizes these data patterns more fully via regression analyses. When workplace flexibility is low, higher investment flexibility is associated with significantly lower planned capital expenditures. In contrast, when workplace flexibility is high (close to one), higher invest-

ment flexibility is associated with significantly higher planned capital expenditures. Showing this *conditional effect* of investment flexibility is unique to our paper. In terms of economic magnitudes, the results in Table 4.3 column (4) imply that for firms with no workplace flexibility, a one standard deviation increase in investment flexibility (0.297) would reduce planned 2020 capital expenditure growth by around 5.9 percentage points ($\approx -0.199 \times 0.297$). For firms with full workplace flexibility, in contrast, a one standard deviation increase in investment flexibility would boost 2020 capital expenditure growth by 16.5 percentage points ($\approx 0.555 \times 0.297$). The magnitudes are somewhat larger in column (5). Figure C.3 provides visualization of the marginal effects of investment flexibility on capital spending plans across different levels of workplace flexibility. Results in columns (1) and (2) of the table further show that the interaction between investment flexibility and workplace flexibility also has some impact on planned employment growth (albeit statistically less significantly so). Since investment flexibility relates to firms' ability to adjust investment spending (it is not directly related to employment), labor and capital need to be complementary for this interaction to affect employment (e.g., low workplace flexibility makes firms use high investment flexibility for reducing employment). Otherwise, this interaction would primarily affect capital spending instead of employment.[14]

Finally, other variables that affect whether firms experience adverse or favorable conditions may also interact with investment flexibility in shaping capital spending decisions. One such variable is financial flexibility. The estimates in column (6) suggest that firms with low financial flexibility (possibly experiencing more adverse conditions) appear to use higher investment flexibility to reduce or delay capital expenditures. In comparison, firms with more financial flexibility (possibly experiencing more favorable conditions) do not use higher investment flexibility to cut capital expenditures, though the interaction term of investment flexibility and financial flexibility is not statistically significant. Column (3) shows similar tendencies when the outcome is employ-

[14] We note that in Table 4.3, given the interaction with investment flexibility, the non-interacted coefficients on workplace flexibility in all columns and the non-interacted coefficient on financial flexibility in columns (3) and (6) reflect the impact of these variables when investment flexibility is set to zero. The interaction terms show that the impact of both workplace and financial flexibility tend to be higher when investment flexibility is positive rather than zero.

ment.[15] Another variable that can affect a firms' prospects is customer demand, and in Table C.8in the Internet Appendix we examine its interactive effects with investment flexibility. This analysis shows that firms with higher customer interaction (hence lower demand and worse conditions during the pandemic) appear more inclined to use higher investment flexibility to reduce or delay capital spending in 2020. The relation has the expected sign but is statistically insignificant in the survey data (Panel A of Table C.8); however, this result is stronger and statistically more significant in Compustat data (Panel B). Overall, how firms use investment flexibility suggests that low workplace flexibility seems to be a key constraining margin during this crisis: workplace flexibility is the strongest determinant of how firms use their investment flexibility to adjust capital spending in 2020.

In our main tests, we use an industry-level measure of investment flexibility, constructed using firm-level data from the March 2019 CFO survey. In order to corroborate our findings in Table 4.3 and provide basis for aggregating investment flexibility to the industry level, we also examine the subsample of firms that responded to both the March 2019 and March 2020 CFO surveys. This allows us to use firm-specific (rather than industry-average) investment flexibility for each firm in this analysis. Table C.3 displays the results: the conditional impact of investment flexibility on investment is present even when we use the firm-level measure and the smaller subsample (47% of the main March 2020 sample).

The findings in this section demonstrate how multiple margins of corporate flexibility affect firms' real decisions during the COVID-19 crisis. The classic margin of financial flexibility plays an important role. Furthermore, we find evidence that workplace and investment flexibility are key factors in shaping firms' hiring and investment decisions for both constrained and unconstrained firms, in line with the predictions of the simple model presented in Section 4.3. As we show below in Section 4.4.2, these results were not observed prior to the COVID-19 crisis.

[15] Table C.7in the Internet Appendix displays specifications containing interactions amongst all pairs of the flexibility measures, showing that the interaction between workplace flexibility and investment flexibility is significantly robust, while other interactions are insignificant.

Other Measures of Flexibility

In addition to our three focal measures of flexibility, we collect data on four other measures that are broadly related to the concept of flexibility. We analyze these measures in this section.

First, we study the role of fixed costs: firms with more fixed costs could also be viewed as being less flexible in changing their expenditures in a crisis. During the COVID-19 outbreak, as revenues decline, firms with high fixed costs may face significant financial pressures. Correspondingly, financial flexibility (such as more cash holdings) can be especially important in helping firms with high fixed costs cover their costs and stay afloat. In contrast, high fixed cost companies with limited financial resources may be forced to lay off workers especially aggressively in order to save cash and cover the fixed costs. Hence, financial flexibility is particularly important when a firm's cost structure is predominantly fixed. In comparison, companies with low fixed costs can reduce production, and both revenues and costs will decrease, so having financial flexibility is less critical for low fixed cost firms.

To measure the prevalence of fixed *vs*. variable costs, we follow prior work (Anderson et al., 2003; Chen et al., 2019) and regress log change in operating costs on log change in sales. The regression coefficient indicates the fraction of variable costs in total costs: if costs are 100% variable, the coefficient would be 1; if costs are 100% fixed, the coefficient would be 0. We use Compustat data to perform this regression for each industry, and apply the estimated variable cost share to firms in our CFO surveys according to their industries. Table 4.4 uses data from the March 2020 CFO survey and shows that planned employment growth during COVID is especially sensitive to financial flexibility among firms with a higher share of costs that are fixed (a lower share of variable costs).[16] In Table C.4, we use Compustat annual employment growth data for 2020 to explore the interaction of fixed costs with financial flexibility. The Compustat-based analysis uses the same industry-level measure of fixed cost share and uses cash holdings and leverage to proxy for financial flexibility; this analysis finds significant evidence that financial flexibility is most

[16] The interactive fixed cost result becomes weaker in column (4) with the inclusion of state fixed effects due to the influence of a small state (District of Columbia, which has four observations) with an industry make-up dominated by investment funds, which have a hard to measure cost structure. Additional details available from the authors upon request.

beneficial when cost structure is predominantly fixed.

Second, we study other dimensions of the workplace environment, including the use of part-time workers and scheduling flexibility. One hypothesis is that if a company can respond more easily by reducing part-time workers or providing more flexible work schedules, it might be better positioned to support more full-time employment. We measure the prevalence of part-time workers and the degree of scheduling flexibility in each industry using the General Social Survey and ATUS, respectively (Mas and Pallais, 2020). We show the relation between these two variables and the growth of full-time employment in 2020 in Table 4.5 columns (1) to (4). We do not observe a significant relation in the data, which suggests that these margins do not appear to be the most critical in the COVID crisis.

Third, we study the effects of labor adjustment costs related to unionization. Our measure of unionization is detailed in Hirsch and Macpherson (2003); it equals the percentage of employees that are union members at the four-digit NAICS level. When unionization is high, firms cannot lay off employees very easily. We study the impact of unionization in Table 4.5 columns (5) and (6). Worker unionization rates do not appear to influence our findings.

Finally, several papers also construct a summary measure of firms' flexibility in changing the scale of operations by examining the variability of operating costs relative to sales (Gu et al., 2019, 2021). This measure is one way to capture adjustment costs in general. We also study this measure in Table 4.5 columns (7) and (8) and do not find a direct relation with firms' employment or investment outcomes in 2020.[17]

4.4.2 A Tale of Two Crises: 2020 vs. 2008

To provide context for our analysis, in this section we characterize and differentiate the impact of a health crisis on firms' decisions from that of other crises, such as those associated with the supply of capital. We do so by comparing corporate decision-making in the COVID-19 crisis to that in the 2008 Financial Crisis. Campello et al. (2010) analyze CFOs' plans for employment and

[17] In untabulated analysis, we repeat Table 4.5 columns (7) and (8) using Compustat data and do not find significant results either, nor do we find significant interactive effects between this measure and workplace or financial flexibility.

investment at the end of 2008 and document the importance of financial flexibility in shaping corporate decisions in the financial crisis. We use the same 2008 CFO survey data to conduct our corporate flexibility analyses, which allows us to compare the effects of flexibility in 2008 *vs.* 2020.

The CFO projections of employment and capital spending growth in the December 2008 survey are for the year 2009. For financial flexibility, we rely on the survey question from December 2008 that asks firms if their operations are affected by difficulties in accessing credit markets. Firms responding "not affected" are classified as having high financial flexibility, while those responding "somewhat affected" and "very affected" are classified as having low financial flexibility.[18] This question focuses primarily on access to credit markets, while the main financial flexibility question in the March 2020 survey captures the ability of firms to access both internal and external funding, as explained in Section 4.2. As a result, the financial flexibility variable in the 2020 survey is broader and may show stronger results for financial flexibility compared to the variable in the 2008 survey. For workplace flexibility and investment flexibility, we use the same industry-level measures as before.

Panel A of Table 4.6 presents the same regression specifications as Table 4.2. Columns (1) and (4) show the results using 2008 data, whereas columns (2) and (5) show the results using the 2020 data. Columns (3) and (6) use the combined sample where we interact workplace flexibility — the distinct central feature of the COVID crisis — and the other flexibility measures with an indicator for the 2020 survey. We find that during *both* the COVID-19 pandemic and the 2008 Financial Crisis, financial flexibility plays a similarly important role in shaping firms' employment and investment plans. However, workplace flexibility is uniquely important for employment plans in the 2020 crisis, while its coefficient in the 2008 data is nearly zero. In analogous fashion, Panel B of Table 4.6 follows the regression specifications in Table 4.3 and shows that firms exploiting their investment flexibility conditional on their workplace flexibility is unique to the 2020 pandemic. Here, too, we find no interactive evidence to suggest that workplace flexibility matters for how firms utilize their investment flexibility in the 2008 crisis.

[18] Accordingly, the group labeled "low financial flexibility" ("high financial flexibility") corresponds to the "constrained" ("unconstrained") group in Campello et al. (2010).

Overall, the comparisons in Table 4.6 highlight that the impact of workplace flexibility is absent in the financial crisis, but has become central since the health crisis. Just as the Global Financial Crisis gave rise to an important body of work on financial constraints, the COVID-19 health crisis may spur critical new research on the transformation of the corporate workplace.

4.4.3 External Validation via Realized Outcomes

Our March 2020 survey provides valuable information about corporate planning in real time as the COVID-19 crisis hit. We subsequently collected more data on companies' realized outcomes from both additional surveys and Compustat to verify the robustness of our findings. Since the subsequent analyses cover a different set of firms, they help confirm that our results on the key drivers of corporate decisions in response to COVID-19 hold in general.

Realized Outcomes based on Subsequent Surveys. In the September 2020 survey, we asked CFOs "For your company, how would you assess the level of the following items (employment, capital expenditures, etc.) compared to their levels before the outbreak of COVID-19?" Figure C.8in the Internet Appendix plots the responses, separately for firms with high *vs.* low workplace flexibility. One can see that as of September 2020, firms with high workplace flexibility were less likely to have experienced reductions in employment. Interestingly, they are not less likely to have cut capital expenditures. These findings are consistent with results in Table 4.2 that high workplace flexibility firms in the March survey anticipated higher employment growth, but did not anticipate higher capital spending in 2020. Correspondingly, the implied physical capital to labor ratio is more likely to have decreased for high workplace flexibility firms, consistent with these firms shifting from capital (in the form of structures and equipment) towards labor. Finally, we also asked firms about their current level of remote work (relative to pre-COVID times). Firms with high workplace flexibility are significantly more likely to have increased remote work (orange *vs.* blue in Panel D of the figure), which aligns with our definition of workplace flexibility. Table C.10in the Internet Appendix verifies these results through ordered logit regressions (coefficients presented in odds ratios).

Realized Outcomes based on Compustat and BLS Data. We also perform external validation

of our survey-based analysis using realized outcomes among Compustat firms. Table 4.7 presents regressions that study annual employment growth and capital expenditure growth in 2020. These annual outcomes map closely to the questions in our CFO surveys (which asked firms about their employment and capital spending plans for the year 2020). We continue to use measures of workplace flexibility and investment flexibility by industry (i.e., these industry-level measures are the same as what we used in previous regressions with survey data). For financial flexibility, nonetheless, we have to rely on proxies in Compustat data. Following other work (Fahlenbrach et al., 2021; Ramelli and Wagner, 2020), we use cash holdings and leverage to proxy for financial flexibility, though we acknowledge that these proxies may only partially measure the concept of financial flexibility. Therefore, in this analysis we interpret cash and leverage as control variables for financial flexibility, more so than as precise measures of financial flexibility.

Columns (1) to (3) of Table 4.7 validate the results in columns (1) to (3) of Table 4.2 — workplace flexibility was an important determinant of employment realizations in 2020. Columns (4) to (6) of Table 4.7 validate our results on the conditional impact of investment flexibility. Similar to what we find in Table 4.3 based on firms' internal plans reflected in the March 2020 survey, realized capital spending growth of Compustat firms confirms a significant interactive effect between workplace flexibility and investment flexibility: firms use their investment flexibility to increase capital expenditures if conditions are favorable and reduce capital expenditures if conditions are unfavorable. Figure C.4 presents a graphical summary of regressions that perform "placebo checks" using Compustat data from previous years. Panel A shows that workplace flexibility did not play a positive role in employment growth prior to 2020. Further, Panel B confirms that neither did the interaction between investment flexibility and workplace flexibility in shaping capital expenditures seem to occur before 2020.[19]

At the industry level, the BLS provides data on employment. In untabulated analyses, we verify that cumulative employment growth since the end of 2019 (through December 2020) is significantly positively correlated with workplace flexibility, with a similar magnitude to what

[19] The 2008 Compustat result in Figure C.4 also confirms the survey-based comparison of the 2008 and 2020 crises in Table 4.6.

we find in Table 4.2. In the rightmost columns of Table C.9in the Internet Appendix, we show that there was no significant relationship between workplace flexibility and employment before 2020.

Corporate Flexibility and Stock Returns. Finally, an influential literature has examined the drivers of stock returns at the onset of the COVID-19 crisis. Our analysis complements this research in that we investigate corporate actions during this period. We examine the determinants of firms' real decisions — both *ex ante* plans via CFO survey data and *ex post* realizations via Compustat and other data — and our findings are generally consistent with results from the stock return literature.

Four main sets of results in the stock return literature relate to our work. First, a number of studies document that firms with more cash and lower leverage (higher financial flexibility) experienced higher stock returns when COVID-19 struck (Fahlenbrach et al., 2021; Ramelli and Wagner, 2020; Ding et al., 2021). Using CFOs' assessments of financial flexibility, we find that such firms had higher planned employment growth and capital expenditures growth. Our analysis of Compustat data further confirms that firms with more cash at the end of 2019 had higher realized employment growth in 2020. Second, firms in industries with more ability to work from home (higher workplace flexibility) also had higher stock returns after COVID-19 hit (Papanikolaou and Schmidt, 2022; Favilukis et al., 2020; Pagano et al., 2021). We find that such firms had significantly higher employment growth in 2020 in both our CFO survey and Compustat datasets. Third, firms in industries facing higher customer interactions (thus potentially lower customer demand) witnessed lower stock returns (Fahlenbrach et al., 2021; Pagano et al., 2021). We find that high customer interaction firms had lower realized employment and capital spending growth in Compustat data, though the latter is not significant. Using CFO survey data, in untabulated analysis, we also find a significant negative relation between customer interactions and planned employment growth in surveys conducted in the second half of 2020, though the relation is insignificant in the March 2020 survey data. Finally, results in Fahlenbrach et al. (2021) suggest that customer interaction and financial flexibility may also interact: in particular, cash appears more important for stock returns when customer interaction is high. In Table C.11in the Internet Appendix, we

test this interaction using real outcomes on employment and capital spending in Compustat data, and find a similar result in terms of coefficient signs (although statistically insignificant).[20]

4.5 The Long-Term Impact of the COVID-19 Crisis

Our results thus far describe the role of corporate flexibility in shaping firms' more immediate responses to the COVID-19 outbreak. In this section, we discuss longer-term implications of the 2020 pandemic.

Long-Term CFO Outlook of Employment and Investment. To understand managers' long-term expectations and the extent to which the COVID-19 crisis can have a lasting impact, we asked in the September 2020 survey "When, if ever, do you expect the level of revenue, employment, capital expenditures, and share of workforce working remotely, to return to where it was before the outbreak of COVID-19?" Because a prior question in this survey asked whether firms' current levels were above or below pre-COVID levels, we are able to ascertain which direction a given firm would need to move in order to return to pre-pandemic activity.

We analyze the recovery of revenue, employment, and capital expenditure intensity to pre-COVID levels. For more relevant insights, we focus on firms that have been negatively affected by COVID-19 (or at least have stayed about the same), which comprise the vast majority of respondents. We visualize the responses in Figure C.5. In Table 4.8, we connect firms' responses with their characteristics using ordered logit regressions. The estimated coefficients are presented in odds ratios and a coefficient below (above) one indicates a faster (delayed) return to normal.

We observe that firms with high workplace flexibility expect employment to recover significantly faster compared to low workplace flexibility firms. Meanwhile, these firms expect a slow recovery in capital spending. This dynamic could be driven by the acceleration of the workplace transformation in light of COVID-19: as companies shift to remote work, the primary types of investment will likely move away from traditional capital expenditures, and possibly towards new forms of investment such as intangibles that facilitate flexible collaboration of the workforce

[20] In unreported results, we find a similar interaction (in terms of signs of coefficients) between financial flexibility and customer interactions in determining employment plans in the March 2020 CFO survey sample, though again insignificant.

(Corrado et al., 2009; Eisfeldt and Papanikolaou, 2013; Haskel and Westlake, 2018; Crouzet and Eberly, 2019). Accordingly, slower growth of capital spending going forward may not necessarily reflect weakness among firms (e.g., tight financial constraints or insufficient aggregate demand) but rather a shift in the nature of work and the nature of investment.

The ongoing transformation of the US corporate workplace is also evident from CFOs' expectations about the staying power of remote work. As of the September 2020 survey, barely any CFOs reported that their firm had decreased their amount of remote work since the COVID crisis began. As shown in Figure C.5, about 50% indicated that the level of remote work would go back to the pre-COVID level by the end of 2021, while 40% thought the level of remote work was unlikely to ever return (see also Bartik et al., 2020; Barrero et al., 2021; Eberly et al., 2021). Table 4.8 further shows that CFOs in industries with higher workplace flexibility are more likely to think that the level of remote work will persist for longer, or is unlikely to ever return to pre-COVID levels.

Automation since the COVID-19 Outbreak. Finally, we study the shift towards automation in response to the pandemic, which may be an important aspect of the lasting impact of COVID-19. A growing literature highlights labor could be displaced by technology as times change (Acemoglu and Restrepo, 2020; Kogan et al., 2021). In the December 2020 survey, we asked CFOs "Has your firm implemented (or does your firm plan to implement) automation to reduce labor since March?" Overall, nearly 40% of the firms responded yes. The push to use automation to replace labor is particularly pronounced among large firms (more than 500 employees), where some 60% responded yes, as shown in Panel A of Figure 4.5.

When we analyze automation via regressions, we control for firm size and customer demand. We also include as a control a proxy for the level of human coordination required in the production in each industry (e.g., teamwork, interpersonal communications); these features are difficult to replicate with an automated workforce (indeed, higher human coordination is associated with a lower automation propensity). We find that firms with low workplace flexibility have a higher propensity to increase labor-reducing automation, as shown in the logit regressions reported in Table 4.9 (see columns (1) and (2)). This result is consistent with our findings above that firms with low workplace flexibility expect a slower recovery of employment, which could be driven in part

by their higher propensity to adopt labor-reducing automation. Intuitively, for low workplace flexibility firms that traditionally relied more on employees performing work onsite, switching to automation can decrease disruptions due to health crises, which has become crucial since the pandemic.

It is noteworthy that the coefficient on workplace flexibility in Table 4.9 remains unchanged in column (2) after we control for industry-level automation penetration from 2004 to 2014 (see Acemoglu and Restrepo, 2020). Indeed, workplace flexibility not being correlated with *prior* automation adoption trends suggests that workplace flexibility does not proxy for inherent technological drivers of automation. In other words, the current push for automation among low workplace flexibility firms is new; not a continuation of prior automation trends. The pandemic experience may have accelerated the shift towards automation, especially among low workplace flexibility firms. While these changes prepare firms to better handle future health crises or other disruptions that would make onsite work difficult, some workers can be displaced. This displacement may have lasting consequences. It is thus important to understand which workers are most exposed to labor-reducing automation, which we investigate next.

For firms that increased automation, we also asked, "Which skill positions were affected by the automation you've implemented or plan to implement to reduce your reliance on labor?" On average, low-skill workers are most affected, as shown in Panel B of Figure 4.5. In particular, firms with low workplace flexibility — which show a stronger propensity to automate in the first place — are more inclined to replace low-skill workers (a similar result can be gleaned from the ordered logit estimations in columns (3) and (4) in Table 4.9). It is also noteworthy in Panel B that about 60% of firms with high workplace flexibility plan to use automation to replace high-skill workers (e.g., back office jobs), conditional on increasing labor-reducing automation.

Taken together, these findings suggest that firms in industries with low workplace flexibility may be prompted to change their production processes and workforce profiles. With higher costs due to health risks, firms in low workplace flexibility sectors may replace human labor with automation. This could possibly contribute to a "robot-led recovery" in these sectors in the long run.

4.6 Conclusion

In early 2020, the US experienced its largest economic dislocation in a decade. The crisis was triggered by an unprecedented emergency of global proportions: the rapid spread of the novel coronavirus. We provide information about corporate decision-making in real time as the COVID-19 crisis hit, as well as firms' planning for both the near term and the long term as the crisis unfolds. We do so via a series of CFO surveys that directly track how firms planned to adjust their operations. We focus, in particular, on how companies use three dimensions of corporate flexibility to adapt to the crisis: financial, workplace, and investment flexibility. We show that in light of the COVID-19 crisis, financial flexibility continues to be a significant determinant of company planning, especially among firms with high fixed costs. Workplace flexibility emerges as an additional critical margin that has both direct effects on employment and interactive effects (via investment flexibility) on investment. Investment flexibility also supports firms' emergency responses; companies facing challenging conditions used investment flexibility to cut capital spending during the crisis, while those facing favorable conditions used investment flexibility to increase spending.

Critically, our analysis suggests that workplace flexibility will shape firms' employment and investment decisions in the years to come. Firms may experience long-term changes in the ways they hire and invest, prompted by COVID-19 and the prominence of workplace flexibility. These transformations require new perspectives for understanding the post-pandemic era. In particular, while we expect financial flexibility to continue to be important, our study shows that workplace and investment flexibility can be central for analyzing firms' decisions going forward. In addition, traditional measures of investment such as capital expenditures are likely to become increasingly incomplete in capturing firms' investment activities. Finally, firms may budget more capital expenditures towards automation, which may affect the size and profile of the workforce. The type of long-term adjustment — whether to support remote work or replace workers via automation — is likely to vary by firm, with a given firm's workplace flexibility a central determinant. While our work provides new, early insight into these important issues, more research is needed.

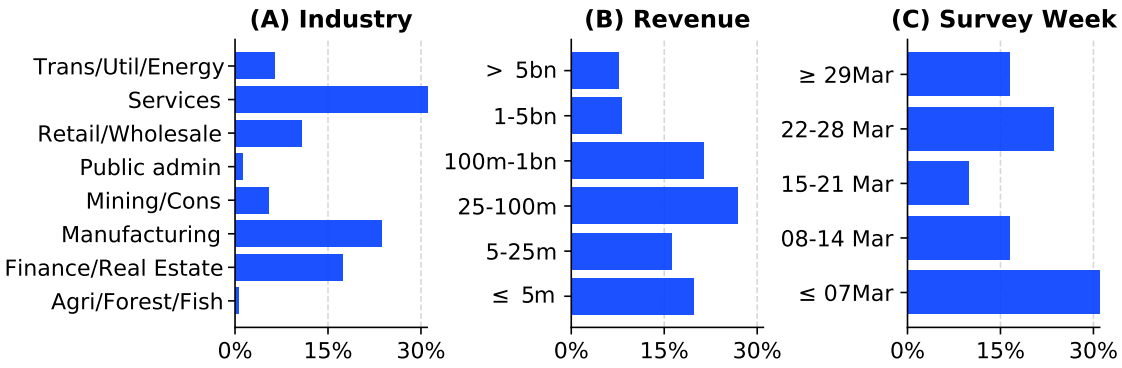

FIGURE 4.1: March 2020 Survey Demographics

This figure shows the composition of firms in the March 2020 survey by industry (Panel A), firm size by 2019 sales revenue (Panel B), and survey week (Panel C). The industries in Panel A are illustrative – the analysis of the paper uses NAICS industry classifications.

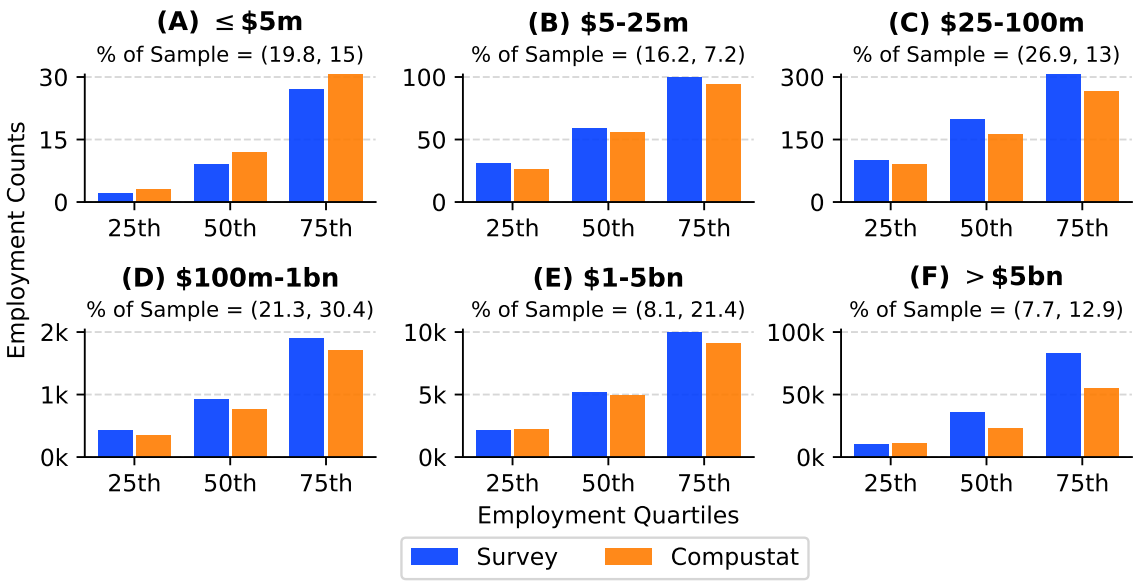

FIGURE 4.2: Comparing March 2020 Survey Firms to Compustat

This figure provides a comparison of firms in the March 2020 CFO survey and Compustat, based on financial information for 2019. It displays the distribution of employee counts across different revenue categories for both survey and Compustat firms, and the percentage of each sample in the revenue category. For example, Panel A displays the employment count quartiles (25th, 50th and 75th percentiles) for both survey and Compustat firms with less than $5m revenue in 2019; and, among these Panel A firms, there are approximately 30 employees at the 75th employment percentile for both the survey sample and Compustat. Below the panel title, the percentage of each sample that falls in the category is displayed (19.8% for the survey sample, 15% for Compustat).

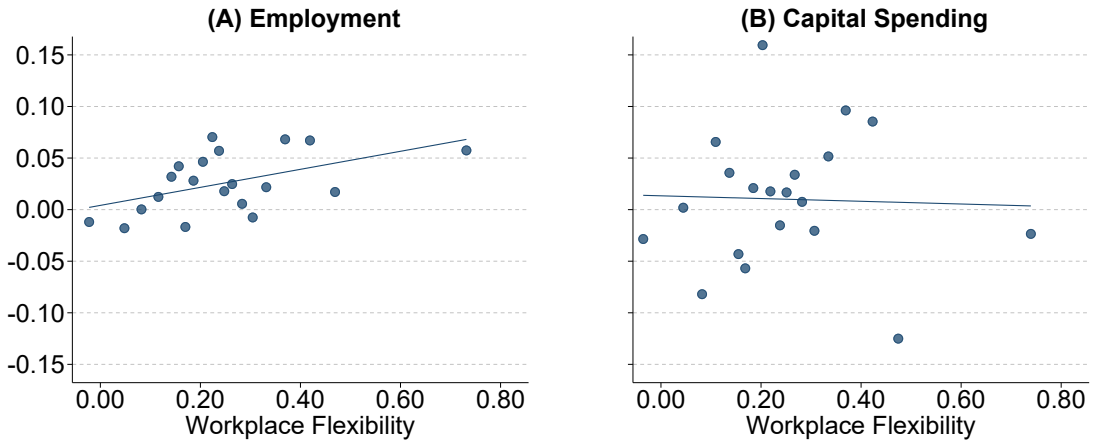

FIGURE 4.3: Direct Impact of Workplace Flexibility on Corporate Plans

Panel A displays a binned scatter plot of CFOs' forecasts of their firms' employment growth in 2020 on workplace flexibility, corresponding to column (3) of Table 4.2, Panel A. Panel B displays the analogous figure for CFO forecasts of their firms' capital spending growth in 2020, corresponding to column (6) of Table 4.2, Panel A. The data come from the March 2020 CFO survey.

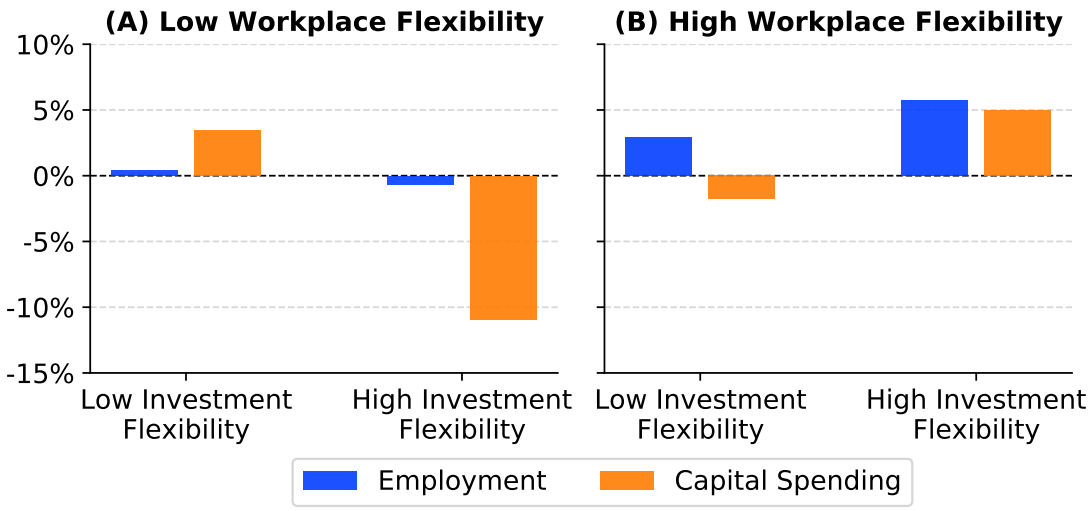

FIGURE 4.4: Impact of Investment Flexibility Conditional on Workplace Flexibility

Panel A displays average CFO forecast of employment and capital spending growth in 2020, for firms with low workplace flexibility (less than or equal to 0.2). Within the panel, average forecasts are shown for firms with investment flexibility below 0.2 (Low) and above 0.2 (High). Panel B displays the analogous figure for firms with high workplace flexibility (above 0.2). The data come from the March 2020 CFO survey.

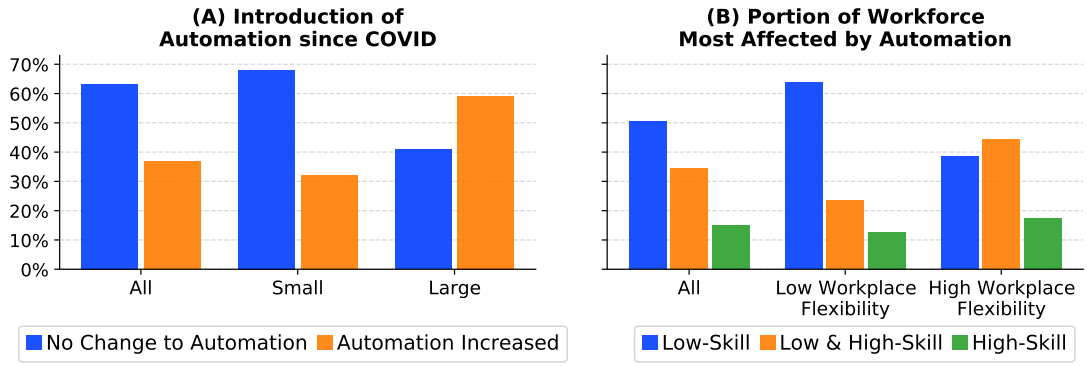

FIGURE 4.5: The Effect of Automation on Employment

Panel A displays the percentage of firms that have increased their use of automation to reduce labor since the onset of the COVID crisis for all firms, small firms and large firms. Large firms have more than 500 employees. For firms that stated their labor-reducing automation had increased, Panel B displays which portion of the workforce will be most affected: low-skill workers only, both low and high-skill workers, or high-skill workers only. Low (high) workplace flexibility is below (above) the 25th (75th) percentile of workplace flexibility within-sample. The data come from the December 2020 CFO survey.

Table 4.1: Descriptive Statistics

	N	Mean	Std dev	25%	Median	75%
CFO Forecast Variables						
Revenue Forecast	501	0.046	0.223	-0.050	0.030	0.100
Employment Forecast	461	0.027	0.175	0	0	0.050
Capital Spending Forecast	453	0.007	0.340	-0.050	0	0.050
Flexibility Variables						
Financial Flexibility	520	0.806				
Workplace Flexibility (ATUS)	451	0.252	0.220	0.064	0.243	0.349
Workplace Flexibility (DN)	454	0.445	0.259	0.225	0.311	0.762
Investment Flexibility	451	0.258	0.297	0	0.200	0.500
Control Variables						
Customer Interactions	451	0.445	0.104	0.369	0.465	0.498
Log # Employees (2019)	520	5.451	2.601	3.807	5.212	6.921

This table presents summary statistics of the main variables for the March 2020 Survey. The number of observations, means, standard deviations, and quartiles are displayed. Detailed variable definitions are given in Appendix C.3.

Table 4.2: Determinants of Employment and Investment Plans

Panel A: Main Specification

	(1)	(2)	(3)	(4)	(5)	(6)
	Employment			Capital Spending		
Financial Flexibility	0.068***	0.071***	0.074***	0.077***	0.081***	0.089***
	(0.017)	(0.018)	(0.019)	(0.027)	(0.026)	(0.027)
Workplace Flexibility	0.100***	0.111***	0.085***	0.032	0.039	-0.007
	(0.029)	(0.030)	(0.024)	(0.043)	(0.041)	(0.064)
Investment Flexibility	0.030	0.017	0.026	-0.030	-0.044	-0.074
	(0.018)	(0.020)	(0.023)	(0.072)	(0.071)	(0.075)
Customer Interactions		0.091**	0.013		0.091	0.055
		(0.041)	(0.080)		(0.117)	(0.265)
Log # Employees (2019)		-0.002	-0.006		-0.004	-0.008
		(0.004)	(0.005)		(0.006)	(0.008)
Observations	405	405	400	397	397	391
R-squared	0.045	0.050	0.224	0.009	0.011	0.159
Week FE			Yes			Yes
State FE			Yes			Yes
NAICS-2 FE			Yes			Yes

Panel B: Alternative Work-from-Home Measure from Dingel and Neiman (2020)

	(1)	(2)	(3)	(4)	(5)	(6)
	Employment			Capital Spending		
Financial Flexibility	0.069***	0.071***	0.075***	0.077***	0.081***	0.091**
	(0.018)	(0.019)	(0.018)	(0.026)	(0.026)	(0.032)
Workplace Flexibility (DN)	0.080***	0.076***	0.093***	-0.006	-0.015	0.001
	(0.019)	(0.022)	(0.030)	(0.047)	(0.051)	(0.060)
Investment Flexibility	0.023	0.018	0.026	-0.026	-0.037	-0.022
	(0.024)	(0.025)	(0.022)	(0.078)	(0.076)	(0.082)
Customer Interactions		0.026	0.055		0.077	0.013
		(0.045)	(0.046)		(0.125)	(0.135)
Log # Employees (2019)		-0.002	-0.006		-0.005	-0.010
		(0.004)	(0.005)		(0.006)	(0.007)
Observations	405	405	400	397	397	391
R-squared	0.042	0.044	0.193	0.009	0.010	0.130
Week FE			Yes			Yes
State FE			Yes			Yes

This table examines the determinants of CFOs' projected growth of employment and capital spending in 2020, using data from the March 2020 CFO survey. The dependent variable is the expected annual growth rate (from the end of 2019 to the end of 2020) of employment (columns (1) to (3)) or capital spending (columns (4) to (6)). In Panel A, Workplace Flexibility comes from ATUS, and is a four-digit NAICS level measure for the percentage of workers that can work from home. In Panel B, Workplace Flexibility (DN) is the work-from-home variable from Dingel and Neiman (2020), measured at the two-digit NAICS level. Financial Flexibility is an indicator taking a value of one if the firm stated they had more financial flexibility than "None" or "A little." Investment Flexibility is a four-digit NAICS level measure for a firm's investment flexibility (with respect to speed of completion). Customer Interactions is a four-digit NAICS level variable that proxies for the intensity of interactions with consumers, which can affect customer demand during the pandemic. Log # Employees (2019) is the natural logarithm of the firm's number of employees at the end of 2019. Detailed variable definitions are in Appendix C.3. Standard errors are clustered at the two-digit NAICS level and displayed in parentheses below the coefficient. ***, **, * denote significance at 1%, 5%, 10%.

Table 4.3: Conditional Impact of Investment Flexibility on Employment and Investment

	(1)	(2)	(3)	(4)	(5)	(6)
	Employment			Capital Spending		
Financial Flexibility	0.071***	0.077***	0.048	0.082***	0.094***	0.059**
	(0.018)	(0.019)	(0.034)	(0.027)	(0.028)	(0.023)
Workplace Flexibility	0.039**	0.046**	0.043**	-0.124	-0.146	-0.150
	(0.017)	(0.017)	(0.017)	(0.073)	(0.100)	(0.100)
Investment Flexibility	-0.037*	-0.025	-0.114*	-0.199**	-0.250***	-0.358***
	(0.021)	(0.027)	(0.057)	(0.074)	(0.077)	(0.087)
Workplace Flex × Investment Flex	0.302***	0.268	0.287	0.754***	0.922***	0.945***
	(0.090)	(0.182)	(0.180)	(0.179)	(0.273)	(0.270)
Financial Flex × Investment Flex			0.104*			0.126
			(0.050)			(0.097)
Observations	405	400	400	397	391	391
R-squared	0.058	0.230	0.234	0.029	0.177	0.179
Controls		Yes	Yes		Yes	Yes
Week FE		Yes	Yes		Yes	Yes
State FE		Yes	Yes		Yes	Yes
NAICS-2 FE		Yes	Yes		Yes	Yes

This table examines the interactive effects of workplace and investment flexibility on firms' employment and capital spending plans. The dependent variable is the CFOs' projected growth rate for employment (columns (1) to (3)) or capital spending (columns (4) to (6)) in 2020, using data from the March 2020 CFO survey. Workplace Flexibility comes from ATUS and is a four-digit NAICS level measure for the percentage of workers that can work from home. Investment Flexibility is a four-digit NAICS level proxy for a firm's investment flexibility (with respect to speed of completion). Financial Flexibility is an indicator taking a value of one if the firm stated they had more financial flexibility than "None" or "A little." Controls are Customer Interactions and Log # Employees (at the end of 2019). Detailed variable definitions are available in Appendix C.3. Standard errors are clustered at the two-digit NAICS level and displayed in parentheses below the coefficient. ***, **, * denote significance at 1%, 5%, 10%.

Table 4.4: Impact of Fixed Costs

	(1)	(2)	(3)	(4)
	Employment			
Financial Flexibility	0.030	0.026	0.031	0.065*
	(0.024)	(0.024)	(0.025)	(0.031)
Fixed Cost Share	-0.112**	-0.122**	-0.114*	-0.055
	(0.043)	(0.045)	(0.054)	(0.057)
Fixed Cost Share × Financial Flex	0.124**	0.136**	0.125**	0.053
	(0.053)	(0.055)	(0.056)	(0.070)
Workplace Flexibility	0.104***	0.118***	0.108***	0.095***
	(0.031)	(0.034)	(0.032)	(0.031)
Investment Flexibility	0.031	0.019	0.020	0.026
	(0.018)	(0.020)	(0.025)	(0.022)
Observations	398	398	398	393
R-squared	0.063	0.069	0.133	0.252
Controls		Yes	Yes	Yes
NAICS-2 FE			Yes	Yes
Week FE			Yes	Yes
State FE				Yes

This table examines the interactive effects of fixed costs and financial flexibility on firms' employment plans. The dependent variable is the expected annual growth rate of employment in 2020, using data from the March 2020 CFO survey. Fixed Cost Share is a four-digit NAICS level variable that proxies for the proportion of a firm's operating costs that are fixed (as opposed to variable). It is measured based on the sensitivity of log changes in operating costs with respect to log changes in sales. Financial Flexibility is an indicator taking a value of one if the firm stated they had more financial flexibility than "None" or "A little." Workplace Flexibility comes from ATUS, and is a four-digit NAICS level measure for the percentage of workers that can work from home. Investment Flexibility is a four-digit NAICS level measure for a firm's investment flexibility (with respect to speed of completion). Controls are Customer Interactions and Log # Employees (at the end of 2019). Detailed variable definitions are in Appendix C.3. Standard errors are clustered at the two-digit NAICS level and displayed in parentheses below the coefficient. ***, **, * denote significance at 1%, 5%, 10%.

Table 4.5: Other Measures of Flexibility

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Employment							
Financial Flexibility	0.069***	0.074***	0.068***	0.074***	0.068***	0.075***	0.068***	0.074***
	(0.018)	(0.019)	(0.017)	(0.019)	(0.017)	(0.019)	(0.017)	(0.020)
Workplace Flexibility	0.102***	0.085***	0.110**	0.092*	0.102***	0.083***	0.100***	0.085***
	(0.025)	(0.023)	(0.045)	(0.049)	(0.031)	(0.025)	(0.029)	(0.024)
Investment Flexibility	0.028	0.029	0.031	0.027	0.031	0.030	0.030	0.027
	(0.021)	(0.023)	(0.019)	(0.023)	(0.019)	(0.023)	(0.019)	(0.024)
Fraction Part-Time	0.019	-0.110						
	(0.098)	(0.135)						
Scheduling Autonomy			-0.016	-0.011				
			(0.035)	(0.053)				
Unionization					0.031	0.128		
					(0.092)	(0.272)		
Scale Inflexibility							0.026	0.033
							(0.394)	(0.339)
Observations	405	400	405	400	405	400	405	400
R-squared	0.045	0.224	0.045	0.224	0.045	0.224	0.045	0.224
Controls		Yes		Yes		Yes		Yes
Week FE		Yes		Yes		Yes		Yes
State FE		Yes		Yes		Yes		Yes
NAICS-2 FE		Yes		Yes		Yes		Yes

This table examines the role of other measures of flexibility in determining firms' employment plans. The dependent variable is the expected annual growth rate of employment in 2020, using data from the March 2020 CFO survey. Financial, Workplace and Investment Flexibility are as defined in previous tables. Fraction Part-Time is constructed from the 2010-2018 General Social Survey files and is the average fraction of part-time to all (part-time + full-time) employees in each four-digit NAICS industry (Mas and Pallais, 2020; Abraham et al., 2018). Scheduling Autonomy is constructed from the 2017-2018 ATUS and is a four-digit NAICS level variable that represents the degree of autonomy workers have in setting their own work schedules (Mas and Pallais, 2017). Unionization is constructed from 2019 BLS data and is the fraction of employees that are members of a union in each four digit NAICS industry (Hirsch and Macpherson, 2003). Scale Inflexibility is the empirical measure of operational inflexibility in Gu et al. (2019) and Gu et al. (2021), which captures the inability for firms to easily adjust the scale of their operations in response to profitability shocks. We take the firm-level measure for the year 2019 and aggregate to the four-digit NAICS level. Controls are Customer Interactions and Log # Employees (at the end of 2019). Detailed variable definitions are in Appendix C.3. Standard errors are clustered at the two-digit NAICS level and displayed in parentheses below the coefficient. ***, **, * denote significance at 1%, 5%, 10%.

Table 4.6: Comparison of 2008 Financial Crisis to 2020 COVID Crisis

Panel A: Determinants of Employment and Investment Plans

	(1)	(2)	(3)	(4)	(5)	(6)
	Employment			Capital Spending		
Financial Flexibility	0.041**	0.066***	0.043**	0.088**	0.087***	0.081**
	(0.017)	(0.022)	(0.017)	(0.032)	(0.024)	(0.036)
Workplace Flexibility	-0.009	0.093***	-0.027	0.167	0.016	0.154
	(0.047)	(0.022)	(0.039)	(0.105)	(0.057)	(0.111)
Investment Flexibility	0.052	0.023	0.029	0.113**	-0.080	0.066
	(0.030)	(0.022)	(0.025)	(0.053)	(0.076)	(0.050)
March 2020 × Financial Flex			0.029			0.010
			(0.027)			(0.044)
March 2020 × Workplace Flex			0.129***			-0.143
			(0.046)			(0.116)
March 2020 × Investment Flex			0.011			-0.129
			(0.031)			(0.082)
Observations	335	400	735	322	391	713
R-squared	0.167	0.188	0.208	0.095	0.124	0.139
Sample	Dec '08	Mar '20	Full	Dec '08	Mar '20	Full
Controls	Yes	Yes	Yes	Yes	Yes	Yes
NAICS-2 FE	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes		Yes	Yes	
State × Survey FE			Yes			Yes

Panel B: Conditional Impact of Investment Flexibility during 2008 and 2020

	(1)	(2)	(3)	(4)	(5)	(6)
	Employment			Capital Spending		
Financial Flexibility	0.041**	0.067***	0.056***	0.081**	0.088***	0.085***
	(0.017)	(0.022)	(0.014)	(0.036)	(0.024)	(0.021)
Workplace Flexibility	-0.008	0.069***	-0.007	0.201	-0.096	0.213**
	(0.051)	(0.014)	(0.053)	(0.123)	(0.067)	(0.098)
Investment Flexibility	0.053**	-0.011	0.039*	0.112	-0.242**	0.078
	(0.021)	(0.028)	(0.021)	(0.074)	(0.085)	(0.079)
Workplace Flex × Investment Flex	-0.004	0.174	-0.041	-0.118	0.827**	-0.022
	(0.179)	(0.152)	(0.145)	(0.196)	(0.297)	(0.228)
March 2020 × Workplace Flex			0.085			-0.317***
			(0.060)			(0.116)
March 2020 × Investment Flex			-0.029			-0.287**
			(0.036)			(0.107)
March 2020 × Workplace Flex × Investment Flex			0.187			0.739**
			(0.196)			(0.338)
Observations	335	400	735	322	391	713
R-squared	0.167	0.191	0.208	0.137	0.139	0.145
Sample	Dec '08	Mar '20	Full	Dec '08	Mar '20	Full
Controls	Yes	Yes	Yes	Yes	Yes	Yes
NAICS-2 FE	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes		Yes	Yes	
State × Survey FE			Yes			Yes

This table examines how different forms of flexibility affect employment and capital spending plans differently in the 2008 and 2020 crises. In Panel A, we run similar tests to Table 4.2, Panel A, and compare the determinants of employment and capital spending across surveys. The dependent variable is the CFOs' projected employment growth in columns (1) to (3), and capital spending growth in columns (4) to (6). In column (1), the sample is the December 2008 CFO survey sample, and the employment growth is for the year 2009. In column (2), the sample is the March 2020 sample, and the employment growth is for the year 2020. In column (3), we combine both surveys and interact our flexibility measures with an indicator variable taking a value of one if the firm is in the March 2020 sample. In column (3), the March 2020 dummy is omitted from the regression as it is colinear with the State × Survey fixed effects. Columns (4) to (6) display similar specifications to columns (1) to (3), with the firm's projected capital spending growth as the dependent variable. In Panel B, we run similar tests to Table 4.3, comparing the effect of the interaction of workplace and investment flexibility on employment and capital spending across surveys. The columns in Panel B follow the same sequence as Panel A. Controls are Customer Interactions (employment counts are not available for the December 2008 survey, thus we exclude the control Log # Employees from this table). Detailed variable definitions are given in Appendix C.3. Standard errors are clustered at the two-digit NAICS level and displayed in parentheses below the coefficient. ***, **, * denote significance at 1%, 5%, 10%.

Table 4.7: Realized Outcomes in Compustat Data

	(1)	(2)	(3)	(4)	(5)	(6)
	Employment Growth			Capital Spending Growth		
Workplace Flexibility	0.155***	0.069**	0.070**	0.121	-0.087	-0.077
	(0.036)	(0.031)	(0.031)	(0.129)	(0.130)	(0.126)
Investment Flexibility	-0.042	-0.049**	-0.030	-0.319***	-0.371***	-0.360***
	(0.026)	(0.024)	(0.021)	(0.099)	(0.095)	(0.093)
Workplace Flex × Investment Flex				0.767***	0.914***	0.897***
				(0.293)	(0.286)	(0.276)
Lagged Leverage		-0.020	-0.026		-0.036	-0.037
		(0.021)	(0.020)		(0.058)	(0.059)
Lagged Cash/Assets		0.217***	0.193***		0.431***	0.430***
		(0.018)	(0.020)		(0.076)	(0.088)
Customer Interactions			-0.219***			-0.123
			(0.075)			(0.317)
Log # Employees (2019)			-0.008***			-0.001
			(0.002)			(0.009)
Observations	4,689	4,689	4,689	4,212	4,212	4,212
R-squared	0.060	0.088	0.093	0.049	0.058	0.059
State FE	Yes	Yes	Yes	Yes	Yes	Yes
NAICS-2 FE	Yes	Yes	Yes	Yes	Yes	Yes

This table examines the effects of workplace and investment flexibility on employment and capital expenditure growth realizations for Compustat firms. Columns (1) to (3) examine the direct effects of workplace and investment flexibility on employment growth realizations using annual Compustat data and columns (4) to (6) examine the conditional effects of investment flexibility on capital expenditure growth realizations for Compustat firms. In columns (1) to (3), the dependent variable is the log change in employment from 2019 to 2020, in columns (4) to (6) the dependent variable is the log change in capital spending from 2019 to 2020. Workplace Flexibility and Investment Flexibility are the same measures used in Tables 4.2 and 4.3. Workplace Flexibility comes from ATUS and is a four-digit NAICS level measure for the percentage of workers that can work from home. Investment Flexibility is a four-digit NAICS level proxy for a firm's investment flexibility (with respect to speed of completion). Lagged Leverage is the firm's 2019 ratio of debt to assets ((DLC + DLTT)/AT), and Lagged Cash/Assets is the firm's 2019 ratio of cash and cash equivalents to assets (CHE/AT). Customer Interactions is a four-digit NAICS level variable that proxies for the intensity of interactions with consumers. Log # Employees is the natural logarithm of the firm's number of employees in the previous fiscal year (2019). To create our sample, we start with all observations for the 2020 fiscal year. We require that a firm have positive assets, non-negative debt, non-missing data for lagged leverage and cash/assets, and a non-missing four-digit NAICS code. Standard errors are clustered at the four-digit NAICS level and displayed in parentheses below the coefficient. ***, **, * denote significance at 1%, 5%, 10%.

Table 4.8: CFO Outlook of Firm Outcomes Returning to Pre-COVID Levels

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Revenue		Employment		Capital Spending		Remote Work	
Workplace Flexibility	0.818	0.842	0.702**	0.772*	1.162*	1.215*	1.233***	1.517***
	(0.128)	(0.122)	(0.166)	(0.132)	(0.082)	(0.100)	(0.068)	(0.119)
Investment Flexibility	1.460***	1.389***	1.221	1.324**	1.233*	1.308**	0.816**	0.873
	(0.118)	(0.124)	(0.132)	(0.116)	(0.118)	(0.115)	(0.103)	(0.130)
Customer Interaction		1.236*		1.101		0.891		1.008
		(0.127)		(0.079)		(0.103)		(0.164)
Log # Employees (2019)		1.261**		2.350***		1.300		3.465***
		(0.113)		(0.154)		(0.178)		(0.154)
Observations	197	197	209	209	210	210	233	233
Pseudo R-squared	0.015	0.025	0.015	0.082	0.007	0.015	0.008	0.124

This table examines how long firms expect the changes brought on by COVID-19 to last. Data are from the September CFO survey. This survey asked CFOs:

When, if ever, do you expect your level of Revenue, Employment, Capital Expenditure (Willingness to Spend on Structures and Equipment), Remote Work to return to where it was before the outbreak of COVID-19?
{0 = No Change, 1 = 2020, 2 = 2021, 3 = 2022, 4 = 2023 or later, 5 = Unlikely to return}

Revenue, Employment and Remote Work refer to the level of the variable. Capital Spending refers to "willingness to spend on structures and equipment." In order to capture how long the negative effects of COVID-19 will last, in columns (1) to (6), we limit the sample to firms that stated their level of the relevant variable (e.g. Revenue in columns (1) and (2)) was the same as or lower than its pre-COVID level. In columns (7) and (8), we limit the sample to firms that stated their level of remote work was the same as or higher than its pre-COVID level. As the dependent variable in each specification has multiple categories, each column presents results from an ordered logit regression, and coefficients displayed are odds ratios (an odds ratio less (greater) than one indicates a decrease (increase)). Workplace Flexibility, Investment Flexibility, Customer Interactions and Log # Employees (at the end of 2019) are standardized to unit variance. Thus, the odds ratios display the proportional change in the odds of observing a higher response from a standard deviation change in the relevant variable. Detailed variable definitions are in Appendix C.3. Standard errors are clustered at the two-digit NAICS level and displayed in parentheses below the coefficient. ***, **, * denote significance at 1%, 5%, 10%.

Table 4.9: Adoption and Use of Automation in Response to COVID-19

	(1)	(2)	(3)	(4)
	Increase in Automation {0 = No, 1 = Yes}		Effect of Automation on Skill {0 = Low, 1 = Low & High, 2 = High}	
Workplace Flexibility	0.758**	0.760**	1.690**	1.605**
	(0.127)	(0.127)	(0.208)	(0.214)
Investment Flexibility	1.134	1.138	0.951	0.915
	(0.155)	(0.155)	(0.220)	(0.220)
Customer Interactions	1.117	1.136	1.352	1.191
	(0.165)	(0.175)	(0.289)	(0.313)
Log # Employees (2019)	1.953***	1.953***	0.667*	0.686*
	(0.145)	(0.145)	(0.213)	(0.218)
Human Coordination	0.294***	0.292***	3.386**	2.841*
	(0.354)	(0.360)	(0.587)	(0.596)
Industry Automation Adoption		1.062		0.445**
		(0.102)		(0.320)
Observations	277	277	102	102
Pseudo R-squared	0.104	0.105	0.073	0.092

This table examines changes to automation since the onset of COVID-19. Data are from the 2020q4 (November/December) CFO survey. This survey asked CFOs two questions about automation:

Since March, has your business implemented, or do you plan to implement automation or technology to reduce your reliance on labor?
{0 = No, 1 = Yes}

Which skill positions were affected by the automation or technology you've implemented or plan to implement to reduce your reliance on labor?
{0 = Low-Skill Workers, 1 = Low & High-Skill Workers, 2 = High-Skill Workers}

In columns (1) and (2), the dependent variable is the CFO's response concerning automation implementation, as described in the first question above. In columns (3) and (4), the dependent variable is the CFO's response concerning automation's effect on different types of workers, as described in the second question above. Columns (3) and (4) focus only on firms that answered yes to the first question. In columns (1) and (2), results are from a standard logit regression. For columns (3) and (4), as the dependent variable has multiple categories, each column presents results from an ordered logit regression. Coefficients displayed are odds ratios (an odds ratio less (greater) than one indicates a decrease (increase)). Workplace Flexibility, Investment Flexibility, Customer Interactions, Log # Employees (at the end of 2019) and Industry Automation Adoption are standardized to unit variance. Thus, the odds ratios display the proportional change in the odds of observing a higher response from a standard deviation change in the relevant variable (in the case of the binary variable Human Coordination, the difference between firms in industries that require low and high levels of human coordination). The variable Human Coordination is a four-digit NAICS level binary variable that takes a value of one for firms in industries that require a high degree of human coordination in the workplace. The variable Industry Automation Adoption represents robot adoption between 2004 and 2014 in different industries constructed by Acemoglu and Restrepo (2020). Detailed variable definitions are in Appendix C.3. Standard errors are clustered at the two-digit NAICS level and displayed in parentheses below the coefficient. ***, **, * denote significance at 1%, 5%, 10%.

5

Conclusion

This dissertation comprises three essays in corporate finance and governance. The first chapter establishes the impact of Say-on-Pay, regularly occurring approval votes on executive compensation, and clarifies the economic mechanism by which this impact occurs. I show that SOP resembles a costly punishment mechanism: shareholders can punish the Board, but doing so is costly. While failure is relatively rare, providing this vote to shareholders has a large impact on firm value. My results provide insight about shareholder voting more generally, suggesting that high support rates in shareholder votes may not imply ineffective monitoring.

The second capture studies the roles of hurdle rates in project development through their impact on bargaining. Our model highlights a trade-off that occurs from using an inflated hurdle rate: passing up projects with moderately positive net present value versus earning a higher share in negotiations from the smaller set of projects. In equilibrium, using an inflated hurdle rate can preserve or even create value, thus offering an explanation for the wide-spread and long-standing use of inflated hurdle rates. The third chapter studies the direct and interactive effects of corporate flexibility at the onset of the crisis. While different forms of flexibility (financial, workplace) are important, our main results highlight that dimensions of flexibility interact with each other.

Appendix A

Appendix to Chapter 2

A.1 Appendix Figures and Tables

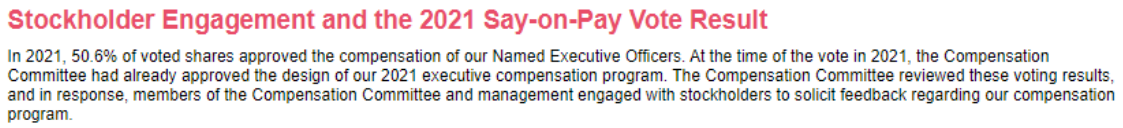

FIGURE A.1: Example of monitoring costs arising from SOP failure — 2021 Netflix SOP

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 (2022).

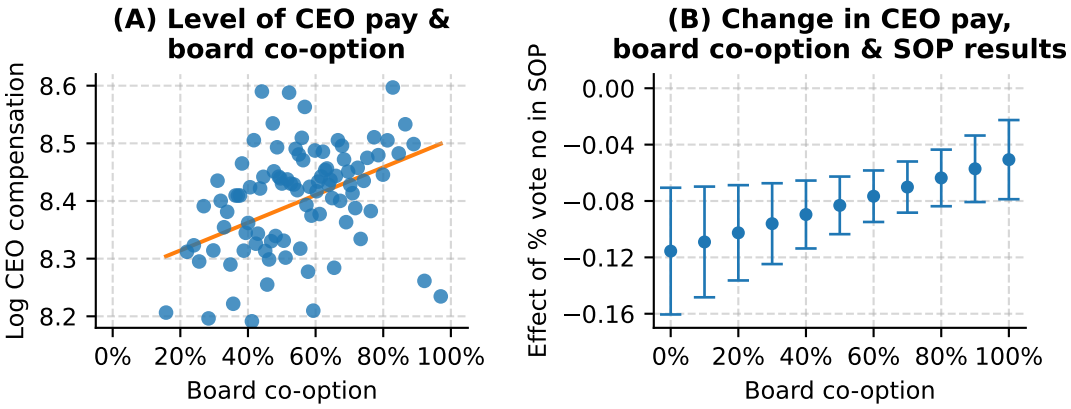

FIGURE A.2: Board capture, CEO pay and SOP results: Illustration of Table 2.3 (Fact 2)

This figure illustrates the results of Table 2.3. Panel A displays a binned scatterplot of the log of CEO compensation on board co-option (as in Coles et al., 2014). Panel B illustrates the effect of SOP disapproval on changes in CEO compensation for varying levels of board co-option. It uses column 6 of Panel B to estimate the relation between changes in CEO pay and the percentage of shareholders voting against the SOP at different levels of board co-option — going from zero to full board co-option, the regression predicts that the effect of a one standard deviation increase in SOP disapproval on the log change in CEO pay increases from -0.12 to -0.05 log points.

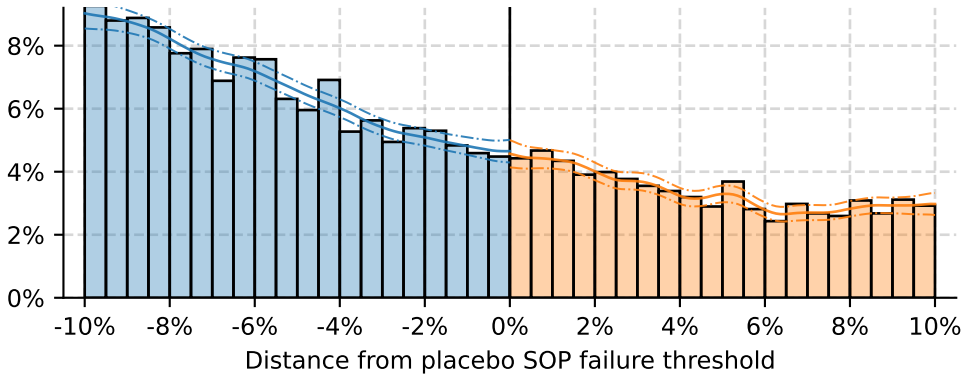

FIGURE A.3: Placebo test of density manipulation of SOP outcomes (robustness for Figure 2.2)

This figure displays the results of placebo testing testing for density manipulation of SOP disapproval at placebo thresholds. Whereas in Figure 2.2, I focus on the publicly accepted important vote failure thresholds of 30% and 50%, in this figure I shift the thresholds to 20% and 40% as a placebo test for density manipulation.

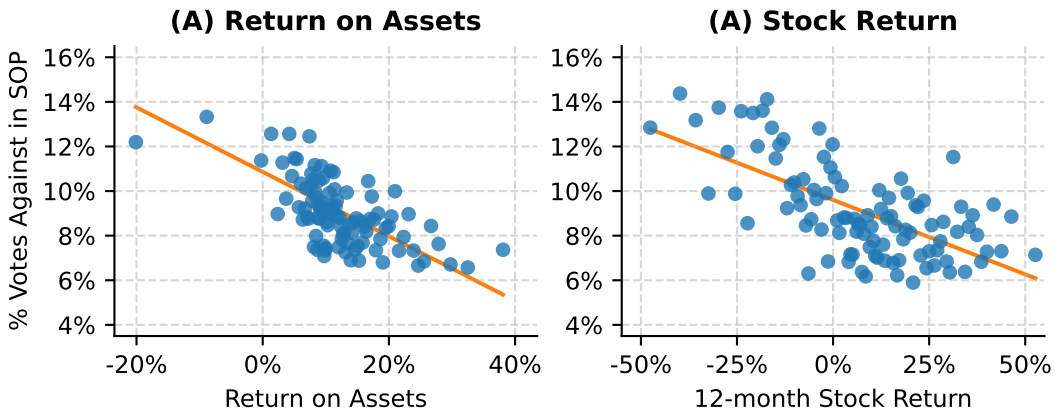

FIGURE A.4: Correlation between SOP % vote against and firm performance (Fact 1)

This figure displays correlations between SOP disapproval rates and measures of firm performance. Both panels display binned scatterplots. Panel A displays results from a regression of SOP disapproval (% vote against in SOP) on the firm's return on assets (ROA). Panel B displays results from a similar regression, where the independent variable is the firm's 12-month stock return. In both panels, I include year and industry fixed effects, and control for log CEO compensation and CEO tenure.

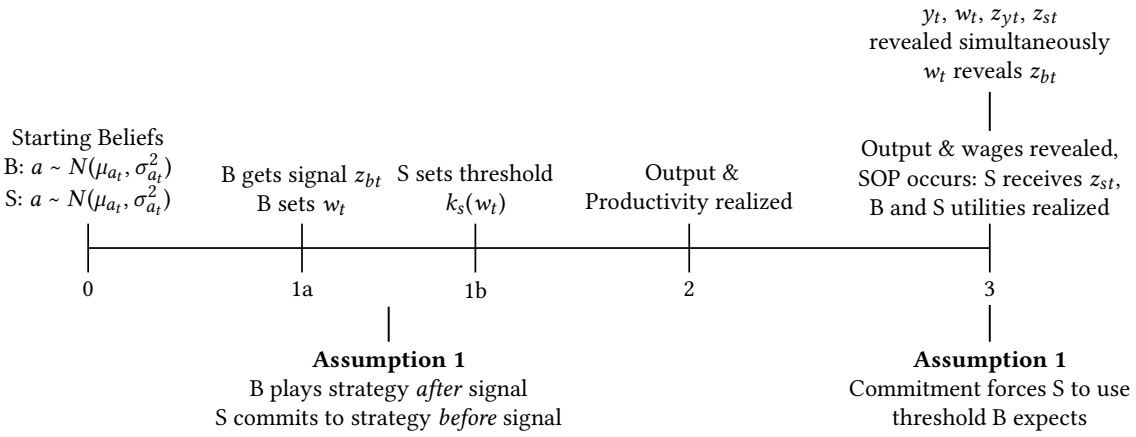

FIGURE A.5: Model timeline with exact time that strategies are played

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 2.3 in the main text for the timeline as it maps to real-world outcomes. See Appendix A.4.4.

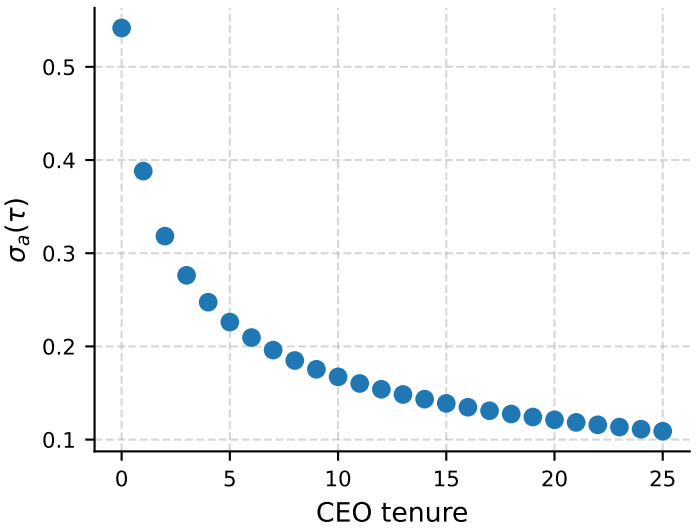

FIGURE A.6: Rate of decline of the variance of beliefs about CEO ability

This figure displays the rate of decline of the variance Board and shareholder beliefs about CEO ability as a function of CEO tenure τ. The figure uses the parameter estimates from Table 2.6. Volatility $\sigma_a(\tau)$ is defined in (D.1)

Table A.1: SOP disapproval and company performance

	(1)	(2)	(3)	(4)	(5)	(6)
	SOP fail $\{0, 1\}$			% vote no in SOP		
Return on assets	-0.027***		-0.024***	-0.020***		-0.018***
	(0.007)		(0.007)	(0.003)		(0.003)
Stock return		-0.016***	-0.014***		-0.012***	-0.011***
		(0.003)	(0.003)		(0.001)	(0.001)
Log firm assets	0.061***			0.030***		
	(0.023)			(0.011)		
Log CEO compensation		0.053***	0.055***		0.027***	0.028***
		(0.009)	(0.010)		(0.005)	(0.006)
Observations	9,864	9,864	9,864	9,864	9,864	9,864
R-squared	0.330	0.329	0.331	0.417	0.417	0.422
CEO tenure FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes

This table displays regressions of SOP vote outcomes (failure and the percentage of shareholders voting against the SOP) on measures of firm performance: the firm's accounting return (return on assets), and the firm's market return (12-month stock return in between SOP votes), in support of Fact 1 and Figure A.4. All covariates are defined in the appendix; 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%.

Table A.2: Changes in CEO compensation following SOP disapproval

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Log change in CEO compensation							
SOP fail $\{0, 1\}$	-0.037***	-0.032***	-0.027**	-0.031**				
	(0.012)	(0.012)	(0.011)	(0.014)				
% vote no in SOP					-0.016***	-0.014***	-0.012***	-0.013***
					(0.003)	(0.003)	(0.003)	(0.004)
Stock return		0.006**	0.005*	0.005		0.005*	0.004	0.004
		(0.003)	(0.003)	(0.003)		(0.003)	(0.003)	(0.003)
Return on assets		-0.003	-0.007	-0.006		-0.004	-0.008	-0.007
		(0.005)	(0.005)	(0.005)		(0.005)	(0.005)	(0.005)
Log firm assets		-0.020**	-0.024**	-0.015		-0.020*	-0.024**	-0.015
		(0.010)	(0.011)	(0.012)		(0.010)	(0.011)	(0.012)
Lagged log CEO compensation				-0.039***				-0.039***
				(0.007)				(0.007)
Observations	7,134	7,134	6,845	4,903	7,134	7,134	6,845	4,903
R-squared	0.176	0.193	0.224	0.226	0.179	0.195	0.226	0.228
Firm FE	Yes	Yes			Yes	Yes		
Year FE		Yes	Yes	Yes		Yes	Yes	Yes
CEO tenure FE		Yes	Yes	Yes		Yes	Yes	Yes
Firm × CEO FE			Yes	Yes			Yes	Yes

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. Stock return 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%.

Table A.3: Model parameters

Description	Notation
Parameters to be estimated	
CEO board capture	λ
Board SOP failure cost	χ_B
Shareholder SOP failure cost	χ_S
Prior average of CEO ability a	μ_0
Prior std dev of CEO ability a	σ_0
Output—CEO wage elasticity	α
Std dev of productivity shock ε_{yt}	σ_y
Std dev of Board signal ε_{bt}	σ_{z_b}
Std dev of Shareholder signal ε_{st}	σ_{z_s}
Scaling factor for estimation	$\log \eta$
CEO wage adjustment cost parameter	c_w
Calibrated parameters	
Board's discount factor	δ_B
Separation probability for CEO of tenure τ	$f(\tau)$

This table displays the parameters of the model, both those to be estimated and those externally calibrated.

Table A.4: Structural estimation results: No shareholder cost

Panel A: Data and model moments

	Description	Notation	Observed	Simulated	t-stat
(1)	Average log output	y_0	7.540	7.539	0.117
(2)	CEO-average output variance	$Var(E_i[\tilde{y}])$	0.880	0.895	-0.961
(3)	Elasticity of output to wage	y_1	0.832	0.847	-0.678
(4)	Output residual variance	$Var(\epsilon^y)$	1.278	1.308	-2.006
(5)	Average log wage when SOP passes	b_0	-0.140	-0.149	3.101
(6)	Change in log wage when SOP fails	b_1	0.042	0.057	-0.991
(7)	Log wage persistence	b_2	0.794	0.916	-12.502
(8)	Wage residual variance	$Var\left(\epsilon^b\right)$	0.059	0.074	-5.981
(9)	SOP failure rate	s_0	0.070	0.046	8.106
(10)	SOP failure—wage sensitivity	s_1	0.061	0.062	-0.076
(11)	SOP failure—output residual sensitivity	s_2	-0.019	-0.020	0.347
(12)	SOP failure residual variance	$Var(\epsilon^s)$	0.064	0.052	4.572
(13)	Bunching estimator	$\boldsymbol{B}$	0.121	0.051	56.535
		χ^2 (p-val)	380.120 (0.000)		

Panel B: Parameter estimates

Description	Notation	Value
Parameters		
CEO board capture	λ	0.092 (0.000)
Board SOP failure cost	χ_B	0.697 (0.001)
Shareholder SOP failure cost	χ_S	—
Prior average of CEO ability	μ_0	0
Prior std dev of CEO ability	σ_0	0.423 (0.000)
Output—CEO wage elasticity	α	0.521 (0.000)
Std dev of productivity shock	σ_y	1.087 (0.001)
Std dev of Board signal	σ_{z_b}	2.430 (0.002)
Std dev of Shareholder signal	σ_{z_s}	0.561 (0.001)
Scaling factor	$\log \eta$	7.612 (0.002)
CEO wage adjustment cost	c_w	3.180 (0.013)
SOP failure costs		
Board SOP cost (% average value)	χ_B / V_0	7.15% (0.01%)
Shareholder SOP cost (% average value)	χ_S / V_0	—

This table contains outcomes of estimation where I remove the shareholder cost to SOP failure as a model parameter. Panel A displays the model fit. In the last row of the panel, I display the result of the χ^2 test of overidentifying restrictions. Panel B displays estimates of parameters.

Table A.5: Parameter estimates and model fit for subsample splits

Panel A.1: Low co-option

Parameters		**Model fit**			
Notation	**Value**	**Notation**	**Observed**	**Simulated**	***t*-stat**
λ	0.357 (0.008)	y_0	7.698	7.698	0.000
χ_B	0.408 (0.026)	y_1	0.939	0.972	-0.818
χ_S	0.041 (0.019)	b_0	-0.141	-0.137	-1.127
μ_0	0.000 (0.000)	b_1	0.044	0.055	-0.445
σ_0	0.500 (0.018)	b_2	0.840	0.860	-1.415
α	0.428 (0.012)	s_0	0.055	0.050	1.207
σ_y	1.045 (0.006)	s_1	0.045	0.036	0.732
σ_{z_b}	0.730 (0.010)	s_2	-0.009	-0.030	4.252
σ_{z_s}	1.075 (0.017)	$Var(E_i[\tilde{y}])$	0.913	0.914	-0.036
$\log \eta$	7.753 (0.018)	$Var(\epsilon^y)$	1.266	1.228	1.493
c_w	4.017 (0.175)	$Var(\epsilon^s)$	0.051	0.057	-1.229
χ_B / V_0	4.17% (0.18%)	$Var\left(\epsilon^b\right)$	0.049	0.061	-3.265
χ_S / V_0	0.42% (0.19%)	$\boldsymbol{B}$	0.046	0.046	-0.160
		χ^2 (p-val)	2.329 (0.507)		

Panel A.2: High co-option

Parameters		**Model fit**			
Notation	**Value**	**Notation**	**Observed**	**Simulated**	***t*-stat**
λ	0.632 (0.007)	y_0	7.729	7.729	0.001
χ_B	0.277 (0.000)	y_1	0.753	0.667	2.577
χ_S	0.178 (0.090)	b_0	-0.110	-0.118	1.708
μ_0	0.000 (0.000)	b_1	0.029	0.052	-1.099
σ_0	0.356 (0.002)	b_2	0.762	0.884	-8.242
α	0.316 (0.007)	s_0	0.080	0.054	5.294
σ_y	0.918 (0.013)	s_1	0.079	0.058	1.607
σ_{z_b}	0.984 (0.021)	s_2	-0.046	-0.032	-2.697
σ_{z_s}	2.223 (0.154)	$Var(E_i[\tilde{y}])$	0.488	0.474	0.524
$\log \eta$	7.763 (0.017)	$Var(\epsilon^y)$	0.967	0.933	1.337
c_w	3.230 (0.123)	$Var(\epsilon^s)$	0.071	0.072	-0.219
χ_B / V_0	3.39% (0.09%)	$Var\left(\epsilon^b\right)$	0.064	0.049	3.666
χ_S / V_0	2.18% (1.07%)	$\boldsymbol{B}$	0.118	0.118	0.251
		χ^2 (p-val)	468.217 (0.000)		

Panel B.1: Low top five institutional ownership

Parameters		Model fit			
Notation	**Value**	**Notation**	**Observed**	**Simulated**	***t*-stat**
λ	0.537 (0.004)	y_0	8.071	8.071	0.000
χ_B	0.401 (0.005)	y_1	1.017	1.041	-0.302
χ_S	0.066 (0.026)	b_0	-0.051	-0.046	-1.190
μ_0	0.000 (0.000)	b_1	0.043	0.052	-0.395
σ_0	0.609 (0.047)	b_2	0.818	0.862	-2.683
α	0.278 (0.007)	s_0	0.062	0.072	-1.805
σ_y	1.193 (0.022)	s_1	0.079	0.066	0.706
σ_{z_b}	1.061 (0.080)	s_2	-0.030	-0.032	0.152
σ_{z_s}	0.636 (0.056)	$Var(E_i[\tilde{y}])$	1.061	1.058	0.033
$\log \eta$	8.080 (0.039)	$Var(\epsilon^y)$	1.615	1.618	-0.024
c_w	5.667 (0.312)	$Var(\epsilon^s)$	0.057	0.056	0.018
χ_B / V_0	2.63% (0.20%)	$Var\left(\epsilon^b\right)$	0.048	0.046	0.108
χ_S / V_0	0.43% (0.16%)	$\boldsymbol{B}$	0.034	0.034	-0.006
		χ^2 (p-val)	1.865 (0.601)		

Panel B.2: High top five institutional ownership

Parameters		Model fit			
Notation	**Value**	**Notation**	**Observed**	**Simulated**	***t*-stat**
λ	0.514 (0.024)	y_0	7.092	7.092	0.000
χ_B	0.462 (0.014)	y_1	0.767	0.739	0.681
χ_S	0.187 (0.042)	b_0	-0.228	-0.225	-0.669
μ_0	0.000 (0.000)	b_1	0.055	0.073	-0.911
σ_0	0.434 (0.008)	b_2	0.757	0.821	-4.008
α	0.368 (0.022)	s_0	0.078	0.023	12.102
σ_y	0.937 (0.021)	s_1	0.089	0.023	4.335
σ_{z_b}	0.756 (0.053)	s_2	-0.016	-0.016	0.051
σ_{z_s}	2.248 (0.130)	$Var(E_i[\tilde{y}])$	0.780	0.663	5.010
$\log \eta$	7.172 (0.017)	$Var(\epsilon^y)$	0.965	1.004	-1.590
c_w	4.121 (0.116)	$Var(\epsilon^s)$	0.070	0.063	2.035
χ_B / V_0	5.63% (0.38%)	$Var\left(\epsilon^b\right)$	0.054	0.091	-8.386
χ_S / V_0	2.28% (0.61%)	$\boldsymbol{B}$	0.140	0.139	1.088
		χ^2 (p-val)	36.575 (0.000)		

Panel C.1: Low blockholder concentration

Parameters		Model fit			
Notation	**Value**	**Notation**	**Observed**	**Simulated**	***t*-stat**
λ	0.590 (0.063)	y_0	8.310	8.310	0.000
χ_B	0.446 (0.021)	y_1	0.752	0.816	-0.837
χ_S	0.091 (0.085)	b_0	-0.012	-0.005	-1.377
μ_0	0.000 (0.000)	b_1	0.036	0.056	-0.964
σ_0	0.579 (0.015)	b_2	0.812	0.819	-0.290
α	0.309 (0.007)	s_0	0.069	0.081	-1.793
σ_y	1.078 (0.021)	s_1	0.077	0.042	1.862
σ_{z_b}	2.848 (0.084)	s_2	-0.019	-0.035	0.796
σ_{z_s}	0.512 (0.033)	$Var(E_i[\tilde{y}])$	0.958	0.926	0.391
$\log \eta$	8.307 (0.038)	$Var(\epsilon^y)$	1.401	1.392	0.099
c_w	5.911 (0.125)	$Var(\epsilon^s)$	0.063	0.076	-2.303
χ_B / V_0	3.69% (0.33%)	$Var\left(\epsilon^b\right)$	0.046	0.086	-1.191
χ_S / V_0	0.75% (0.74%)	$\boldsymbol{B}$	0.046	0.045	0.014
		χ^2 (p-val)	6.769 (0.080)		

Panel C.2: High blockholder concentration

Parameters		Model fit			
Notation	**Value**	**Notation**	**Observed**	**Simulated**	***t*-stat**
λ	0.582 (0.061)	y_0	6.848	6.848	-0.000
χ_B	0.789 (0.043)	y_1	0.858	0.890	-0.798
χ_S	0.365 (0.081)	b_0	-0.267	-0.279	2.987
μ_0	0.000 (0.000)	b_1	0.062	0.095	-1.521
σ_0	0.618 (0.012)	b_2	0.751	0.605	9.470
α	0.328 (0.005)	s_0	0.071	0.047	5.568
σ_y	0.930 (0.024)	s_1	0.090	0.036	3.648
σ_{z_b}	0.318 (0.029)	s_2	-0.030	-0.029	-0.139
σ_{z_s}	3.698 (0.731)	$Var(E_i[\tilde{y}])$	0.836	0.803	0.856
$\log \eta$	6.935 (0.017)	$Var(\epsilon^y)$	1.183	1.175	0.170
c_w	4.655 (0.579)	$Var(\epsilon^s)$	0.064	0.075	-3.308
χ_B / V_0	7.99% (0.05%)	$Var\left(\epsilon^b\right)$	0.056	0.067	-2.285
χ_S / V_0	3.70% (0.64%)	$\boldsymbol{B}$	0.131	0.120	5.984
		χ^2 (p-val)	42.052 (0.000)		

This table displays the parameter estimates with standard errors and model fit for each subsample split presented in Table 2.8. For brevity, I omit descriptions of parameters and moments from the table.

Table A.6: Estimates of the shareholder cost to SOP failure using bunching

Interval $[p^-, p^+]$	(1) ±1%	(2) ±1.5%	(3) ±2%	(4) ±2.5%	(5) ±3%
Bunching **B**	16.61***	11.18***	10.09***	9.60***	7.72***
	(3.16)	(2.76)	(3.16)	(3.03)	(2.51)
Bunching range $\overline{p} - \underline{p}$	7.13***	7.31***	9.22***	10.81***	10.64***
	(1.35)	(1.74)	(2.75)	(3.27)	(3.29)
SOP failure cost $k(\overline{p})$	0.85***	0.89**	1.42*	1.95**	1.89*
	(0.33)	(0.42)	(0.78)	(0.99)	(0.99)
Observations	3,938	3,938	3,938	3,938	3,938
Bootstrap replications	500	500	500	500	500

This table displays estimates of bunching in SOP votes, the bunching range that arises from the bunching mass, and the SOP failure cost that arises from the bunching range. In each column, I specify a region around the punching threshold of $\underline{p} = 30\%$, such that $\mathbf{E} = [p_-, p^+] = \left[\underline{p} - e, \underline{p} + e\right]$, for $e \in \{1\%, 1.5\%, 2\%, 2.5\%, 3\%\}$. In the first column, I display the bunching mass (in percent), as estimated via (E.1) in Section XX. I use a 5th degree polynomial to estimate the counterfactual distribution. The second row displays the estimated bunching range arising from the bunching mass (eq. C.4). The third row displays the estimated SOP failure cost $k(\overline{p})$ (eq. C.3). Bootstrapped standard errors, based on 500 replications, clustered by year to account for time trends in SOP votes, are displayed below the estimates. ***, **, * denote significance at 1%, 5%, 10%.

Table A.7: Moment targeting exercise

	Auxiliary model outcome	Description	Target
(1)	$\hat{y}_0$	Average log output	$\log \eta$
(2)	$\hat{y}_1$	Elasticity of output to wage	α
(3)	$\widehat{Var}(\epsilon^y)$	Output residual variance	σ_y
(4)	$\widehat{Var}(E[\ln y - \hat{y}_1 \ln w])$	CEO-average output variance	σ_0
(5)	$\hat{b}_0$	Average log wage when SOP passes	λ
(6)	$\hat{b}_1$	Difference in log wage when SOP fails	χ_B
(7)	$\hat{b}_2$	Persistence in log wages	c_w
(8)	$\widehat{Var}\left(\epsilon^b\right)$	Wage regression residual variance	σ_{z_b}
(9)	$\hat{s}_0$	Observed SOP failure rate	λ, χ_S
(10)	$\hat{s}_1$	Sensitivity of SOP failure to log wage	χ_B, χ_S
(11)	$\hat{s}_2$	Sensitivity of SOP failure to output shock	σ_{z_s}, σ_y
(12)	$\widehat{Var}(\epsilon^s)$	SOP fail regression residual variance	σ_{z_s}
(13)	$\widehat{\mathbf{B}}$	Bunching estimator (E.1)	χ_B, χ_S

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

Table A.8: Counterfactual analysis: Eliminating the Shareholder cost to SOP failure

	Mean	**25%**	**50%**	**75%**
SOP failure rate	41.68%			
Percent change in				
Wages	-4.74%	-8.64%	-2.97%	0.00%
Firm value	+0.47%	+0.18%	+0.22%	+0.53%

This table displays a counterfactual related to changing the Shareholder costs to SOP failure. I display the counterfactual in which I change χ_S to 0. The first row displays the counterfactual SOP failure rate. Rows 2 and 3 display the counterfactual percentage change in wages and firm value. To compute these changes, 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. I then display the average percentage change, along with the 25th, 50th and 75th percentiles.

A.2 Institutional Details of Say-on-Pay

The 2010 Dodd-Frank Wall Street Reform and Consumer Protection Act, commonly referred to as Dodd-Frank, made SOP compulsory at all US firms from 2011.[1] In the US, SOP is a non-binding vote that must occur at least once every three years on the level and structure of executive compensation. Though SOPs are required only every 3 years, in practice nearly all S&P1500 firms, the main sample, hold the SOP every year; I will commonly refer to SOP being "annual" in the paper, and the model time will be annual. While the vote itself is non-binding, in spirit a low level of shareholder approval for the SOP is likely to lead to tangible changes in the CEO's compensation contract the next year (see Section 2.2.2 and, e.g., Balsam et al., 2016).

By rule, the SOP vote must cover all executive compensation disclosed pursuant to Item 402 of Regulation S-K. This includes the Compensation Discussion and Analysis (CD&A) section of the proxy statement, which is designed to put into perspective the level of executive pay, its structure (e.g., cash *vs.* stock options) and provide a clear narrative of *why* executives received such pay (Dalton and Dalton, 2008).

A key aspect of SOPs are that they are *backward-looking* and *reactive*. From Novick (2019), "Say-on-pay votes ask shareholders to opine retrospectively on the compensation of named executives that is disclosed in the proxy statement, rather than on the company's compensation program going forward." SOPs in the US are clearly a non-binding confidence vote in the Board's choice of CEO wage for the previous fiscal year, and not a vote in Shareholders' confidence about the next year's compensation contract. This backward-looking aspect will inform the timing of the model. In particular, the practice of the compensation committee setting the wage before and shareholders voting on the wage will play a key role in how the model structures the strategies of the Board and Shareholders.

What constitutes SOP failure? By nature in the US, they are non-binding votes, so there is no threshold at which the compensation committee must make a tangible change. Nevertheless there are three important thresholds for the vote. The most important is 70% support. If sup-

[1] Technically, firms with revenue less than $1billion did not have to implement SOP until 2013.

port falls below this, Institutional Shareholder Services (ISS) will publicly push the compensation committee and firm more generally to make changes to compensation policy and engage with shareholders (ISS, 2022). Further, 80% (the threshold at which Glass-Lewis will pursue the Board) and 50% (the classic simple majority) also represent important thresholds.

Upon SOP failure, the firm's compensation committee will often reach out to the firm's large stockholders with a proposed change to the firm's compensation policy in future years, or simply to discuss how these large shareholders feel about the firm's compensation policy. For example, Figure A.1 displays an example of SOP failure, where the Netflix Board reach out to large shareholders to discuss compensation policy. The model is silent on the repeated nature of this interaction. In the model, there is no retrospective action if SOP fails, rather the Board and shareholders pay the utility cost from failure, and the model moves to the next period. I am not interested in perfectly modeling the game between the Board and shareholders, I merely look to put sensible structure on the data to be able to identify the key parameters driving incentives in this setting.

A.3 Empirical Appendix

A.3.1 Estimates of the Shareholder Cost Directly Via Bunching

This appendix presents an alternate structural model which uncovers a reduced-form object which maps to a SOP failure cost, in support of Fact 4 in Section 2.2.2 and the estimates from the structural model presented in Sections 2.3 and 2.4. The magnitudes of the SOP failure cost identified via the model in this Appendix are consistent with my main estimates. However, the key issue is that I cannot separately identify the Board and Shareholder costs to SOP failure here; in my main estimation, I use variation coming from endogenous patterns in CEO pay and SOP vote outcomes to separately identify the Board and Shareholder SOP failure costs.

A simple structural model of costs to SOP failure

Each firm's shareholder base is composed of a group of shareholders, who hold a shareholder approval vote on the Board's pay policy (a SOP vote). The outcome of the vote is $p \in [0, 1]$; $p = 0$ means perfect support for the Board. I abstract away from strategic considerations: the

shareholder base coordinate to pick the p that maximizes value. I assume that each p elicits a unique wage and hence each p has a unique effect on firm value, so there is a unique p^* which maximizes firm value. There is an exogenously decided threshold $\underline{p}$, if $p \geq \underline{p}$ then the SOP fails, i.e. $\underline{p} = 30\%$. The shareholders choose p relative to the undistorted maximum p^* to minimize the cost

$$\min_{p} \quad C(p - p^*) + K(p^*) \times \mathbf{1}\left[p \geq \underline{p}\right] \tag{C.1}$$

The first term captures the cost that shareholders incur if the vote choice p is different from the choice p^* that maximizes firm value (the undistorted vote), i.e. the cost incurred from the CEO receiving wage $w(p)$ instead of $w(p^*)$. C has functional form

$$C(p - p^*) = \frac{1}{2} V(p^*) p^* (1 - p / p^*)^2 \tag{C.2}$$

where $V(p^*)$ is firm value under the undistorted vote. Thus, $C(p - p^*)/V(p^*)$ represents the cost from choosing vote p (when p^* is optimal in the frictionless sense) in percentage terms of firm value. The quadratic cost is standard in the bunching literature (Ewens et al., 2023), and captures in simple form that the cost increases the further the vote outcome is from the undistorted vote which would maximize firm value.

The second term of the payoff function is the cost of SOP failure to shareholders, which is paid if p is above $\underline{p}$. K is a constant, and the discrete jump in cost directly at the SOP failure threshold incentivizes bunching directly at $\underline{p}$. There is a marginal vote outcome, $\overline{p}$, for which shareholders are indifferent between between bunching at $\underline{p}$ and issuing vote $\overline{p}$. The indifference of the marginal vote reveals the cost of SOP failure

$$k(\overline{p}) = \frac{K(\overline{p})}{V(\overline{p})} = \frac{C\left(\underline{p} - \overline{p}\right)}{V(\overline{p})} = \frac{1}{2}\overline{p}\left(1 - \underline{p} / \overline{p}\right)^2 \tag{C.3}$$

Thus, knowing $\overline{p}$ reveals the shareholder cost, relative to firm value, of failing the SOP. To estimate $\overline{p}$, I can use the observed degree of bunching. In the absence of SOP failure cost, the distribution of votes would be smooth around the threshold $\underline{p}$. In the presence of the SOP failure

cost, blockholders that fall in $\left[\underline{p}, \overline{p}\right]$ bunch at $\underline{p}$ as the SOP failure cost exceeds the benefit of being at the undistorted maximum. Let f_0 denote the density function of the counterfactual distribution of votes as if there were no SOP failure cost. As Saez (2010); Alvero and Xiao (2023); Ewens et al. (2023) point out, the degree of observed bunching is the fraction of agents that fall in $[\underline{p}, \overline{p}]$ under the counterfactual distribution, or

$$\mathbf{B} = \int_{\underline{p}}^{\overline{p}} f_0(x)\, dx = \left(\overline{p} - \underline{p}\right)\overline{f_0}$$

where $\overline{f_0}$ is the average density within the bunching region.[2] Further, bunching can also be written as

$$\mathbf{B} = \int_{p_-}^{p^+} f(x) - f_0(x)\, dx$$

where f is the density function of the observed distribution of votes, and $[p_-, p^+]$ is a narrow region around the threshold $\underline{p}$. The intuition is that this definition of bunching captures the amount of mass that is shifted from $\overline{p}$ to $\underline{p}$, relative to the counterfactual distribution. Note that this definition of bunching is a reduced form object in the data, and we can estimate $\overline{f_0}$ by focusing on the average counterfactual density in $[\underline{p}, p^+]$. Equating the two definitions gives

$$\left(\overline{p} - \underline{p}\right) = \frac{\mathbf{B}}{\overline{f_0}} \tag{C.4}$$

which is a reduced-form object idenfiable in the data (Alvero and Xiao, 2023), and identifies the bunching range, which can be used to estimate the SOP failure cost.

To proceed, I estimate $\hat{\mathbf{B}}$ as in Section 2.4, by choosing a region $\mathbf{E}$ close to the threshold, and estimating the counterfactual distribution function as a polynomial of observed counts $\notin \mathbf{E}$. $\hat{\overline{f_0}}$ is the average density in $\left[\underline{p}, p^+\right]$ from this counterfactual distribution. The ratio of these two objects gives the bunching range, which is used to estimate (C.3).

Table A.6 displays the results. In each column, I set a different region $\mathbf{E}$, from $\left[\underline{p} - 1\%, \underline{p} + 1\%\right]$ to $\left[\underline{p} - 3\%, \underline{p} + 3\%\right]$ increasing by 0.5 percentage points in each column. The choices of $\mathbf{E}$ are motivated by Figure 2.2, which shows that this is where the bunching is most apparent. To estimate

[2] That is, $\overline{f_0} = \left(\overline{p} - \underline{p}\right)^{-1} \int_{\underline{p}}^{\overline{p}} f_0(x)\, dx$.

standard errors, I do 500 bootstrap iterations for each **E**, and display the error below the estimate. The first row displays the estimated bunching (in percent) for that interval. The degree of bunching ranges from 8% to 17% of observed votes. The second row displays the estimate of the bunching range $\overline{p} - \underline{p}$ (in percentage points). For example, given that $\underline{p} = 30\%$, the first columns shows that the estimated indifference point in the vote is $\overline{p} = 37.13\%$. For each threshold, the degree of bunching and the bunching rage are both statistically significantly different from zero.

The object of interest, the estimated SOP failure cost is displayed in the last row, and ranges from 0.85% to 1.95%. While the magnitude of these costs is identified via the functional form of C, they are entirely consistent with the estimated parameter from my main structural model. However, this simple model cannot disentangle if this is the Shareholder cost or the Board cost. The threshold $\overline{p}$ is a function of the Board's wage best response (which is in part determined by the magnitude of the Board cost) as much as it is the shareholder SOP failure cost.

A.4 Model Appendix

A.4.1 CEO Contract and Board Capture Microfoundation

This appendix presents a simple microfoundation of the connection between CEO effort, the CEO wage and output. It also microfounds the Board's bias. I show that I cannot separately identify CEO effort aversion from the parameter that governs output's curvature in CEO effort; and the CEO's outside option from the Board's explicitly bias towards overpayment.

The CEO is effort-averse and will work for any wage. That is, given effort-aversion parameter $\gamma \in (0, 1)$, the amount of effort the CEO gives for a given w_t is

$$n_t = w_t^{\gamma}$$

This captures, in a reduced form sense, that extracting more CEO effort increases as more effort is needed. Output is decreasing returns-to-scale in CEO effort,

$$y_t = A_t(n_t^*)^{\beta} = A_t w_t^{\gamma\beta}, \quad \beta \in (0, 1)$$

As such, $y_t = A_t w_t^{\gamma\beta}$, and γ and β are not separately identified. I define $\alpha = \gamma\beta \in (0, 1)$ as my

main parameter of interest, which captures the curvature of production to changes in CEO wage, embedding how the wage affects effort.

In the static model with no adjustment costs and no SOP, the Board solves

$$\max_{w_t} \quad A_t w_t^{\alpha} - \left(1 - (\overline{u} + \overline{b})\right)^{1-\alpha} \times w_t$$

where $0 \leq \overline{u} + \overline{b} < 1$ represents *board capture*, or the degree to which the Board wants to overpay the CEO relative to the profit-maximizing wage. Formulating the problem in this manner is convenient, as the bias term becomes a constant in logs. The CEO's outside option is $\overline{u}$, and the Board's explicit bias is $\overline{b}$. This term is raised to $1 - \alpha$ so the exponent drops out in the first-order condition. Solving the Board's problem and taking logs reveals that

$$\log w_t^{\mathrm{B}} = \log\left(\frac{1}{1 - \left(\overline{u} + \overline{b}\right)}\right) + \log\left((\alpha A_t)^{\frac{1}{1-\alpha}}\right) = \log\left(\frac{1}{1-\lambda}\right) + \log\left((\alpha A_t)^{\frac{1}{1-\alpha}}\right)$$

As such, $\overline{u}$ and $\overline{b}$ are not separately identified, so I define λ, $0 \leq \lambda \leq 1$ to capture board capture. The Board's bias factor is $\frac{1}{1-\lambda}$ and describes how much more the Board wants to pay the CEO., and in log terms the bias shows up as a constant.

A.4.2 Microfoundation of Representative Shareholder Assumption

This section microfounds my assumption of a representative shareholder, under the assumption that a diffuse shareholder base receive a correlated private signal of CEO ability (see Section 2.3.2). Informally, when all atomistic shareholders receive a correlated signal and vote with the same threshold strategy, there is informational equivalence in focusing on a representative shareholder that aggregates information across the shareholder base into a single signal and votes with the same threshold strategy. 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, negative signal of CEO ability. While I remain largely silent on the role of proxy advisors in the SOP process, I do acknowledge their importance.

Proposition A.4.1. *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 (2022). 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

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

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

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

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

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

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

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

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 $Pr(SOP_i = 1 \mid \overline{z})$ is anything other than p. Given that $\overline{z}$ completely determines $Pr(SOP_i = 1 \mid \overline{z})$, there is a bijection between $\overline{z}$ and p

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

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

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

The likelihood of observing p is driven by the likelihood of observing $\overline{z}(p)$, scaled by a change-of-variable term $\overline{z}'(p)$. Since $\overline{z}$ is conditionally normal around the type a, I integrate over all types a and then take the likelihood of observing $\overline{z}(p)$ 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 $\overline{z}$, given a, is equal to $\overline{z}(p)$, which in the limit is the

only $\overline{z}$ for which I would see p. This is a scaled product of Gaussians, so the posterior is also normal,

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

where

$$\mu_{ap} = \frac{\sigma^2_{\overline{z}}}{\sigma^2_a + \sigma^2_{\overline{z}}}\mu_a + \frac{\sigma^2_a}{\sigma^2_a + \sigma^2_{\overline{z}}}\overline{z}(p), \qquad \sigma^2_{ap} = \frac{\sigma^2_a \sigma^2_{\overline{z}}}{\sigma^2_a + \sigma^2_{\overline{z}}}.$$

Thus, observing p is informationally equivalent to observing a signal $z_s = \overline{z}(p) = k_s(w) - \sigma_z \Phi^{-1}(p)$, where $z_s = a + \varepsilon_s$, and $\varepsilon_s \sim N(0, \sigma^2_{z_s})$.

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.

A.4.3 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.

Evolution of Beliefs Period to Period

Prop. A.4.2 shows how beliefs change from t to $t+1$. Prop. A.4.3 describes the distribution of next period beliefs given today's beliefs, which is used when the Board calculates their (expected) continuation value.

Proposition A.4.2. *From period t to $t+1$, the variance of beliefs for both the Board and shareholders declines deterministically according to*

$$\sigma^2_a(\tau + 1) = \left[\sigma^{-2}_a(\tau) + \sigma^{-2}_{z_b} + \sigma^{-2}_{z_s} + \sigma^{-2}_y\right]^{-1} \tag{D.1}$$

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

$$\sigma_a^2(\tau) = \sigma_0^2\left[1 + \tau\left(\kappa_{z_b}^{-1} + \kappa_{z_s}^{-1} + \kappa_y^{-1}\right)\right]^{-1} \tag{D.2}$$

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

$$\mu_{at+1} = \sigma_a^2(\tau+1)\left[\frac{\mu_{at}}{\sigma_a^2(\tau)} + \frac{z_{bt}}{\sigma_{z_b}^2} + \frac{z_{st}}{\sigma_{z_s}^2} + \frac{z_{yt}}{\sigma_y^2}\right] \tag{D.3}$$

Proof The formulas are standard results in Bayesian learning (e.g., Pastor and Veronesi, 2009; Taylor, 2010).[3] The Board and shareholder reveal their signals each period. Thus, Board and shareholders share the same beliefs about the variance from period to period.

Proposition A.4.3. *The mean and variance of the mean of $t+1$ CEO beliefs at t are*

$$\begin{aligned} E_t[\mu_{at+1}] &= \mu_{at} \\ Var_t[\mu_{at+1}] &= \sigma_a^2(\tau) - \sigma_a^2(\tau+1) \end{aligned} \tag{D.4}$$

That is,

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

Proof I drop time subscripts for convenience, and use $\cdot'$ to denote next period. Via Prop. A.4.2, 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,

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

[3] See also the internetappendix for Taylor (2010).

which of course must hold because beliefs are a martingale. 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

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

and I am done. Equivalently, I can write this as $\frac{p' - p}{p'p}$. Further, 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

Differences in Board and Shareholder Beliefs Within Period

This subsection explains exactly how Board and shareholder beliefs evolve within each period. As the wage and vote perfectly reveal signals z_b and z_s, the Board and shareholders share the same beliefs at the beginning of any period. Let τ_t be the tenure of the CEO at time t. By Prop. A.4.2, I have that $\sigma_{bt}^2 = \sigma_{st}^2 = \sigma_a^2(\tau)$ from (D.2). At the beginning of the period, let μ_{at} be the beliefs about the mean at the beginning of period t. So, I can describe Board and shareholder beliefs as $\big(\mu_{at}, \sigma_a^2(\tau_t)\big)$ at the beginning of period t.

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

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

$$\mu_{bt|z_b} = \sigma^2_{bt|z_b}\left(\frac{\mu_{at}}{\sigma^2_a(\tau_t)} + \frac{z_{bt}}{\sigma^2_{z_b}}\right) \tag{D.5}$$

$$\sigma^2_{bt|z_b} = \frac{\sigma^2_a(\tau_t)\sigma^2_{z_b}}{\sigma^2_a(\tau_t) + \sigma^2_{z_b}} = \sigma^2_0\left[1 + (\tau_t + 1)\kappa^{-1}_{z_b} + \tau_t(\kappa^{-1}_{z_s} + \kappa^{-1}_y)\right]^{-1} \tag{D.6}$$

   I use $\mu_{bt|z_b}$ and $\sigma_{bt|z_b}$ to follow conventions in the main text. The Board makes their wage decision based upon these beliefs. Before the wage is revealed, shareholders still maintain beliefs $\left(\mu_{at}, \sigma^2_a(\tau_t)\right)$.

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

   When Board and shareholders play the wage-SOP game, their beliefs differ, in that the Board believes $\left(\mu_{bt|z_b}, \sigma_{bt|z_b}\right)$ and shareholders believe $\left(\mu_{at}, \sigma^2_a(\tau_t)\right)$. But shareholders can discern $\left(\mu_{bt|z_b}, \sigma_{bt|z_b}\right)$ for any z_{bt}, which they factor in when choosing their threshold. Shareholders thus factor in what $\mu_{bt|z_b}$ will be when calculating expected wages in their objective function in (2.17).

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

$$Z_{syt} = a + p\varepsilon_{st} + (1 - p)\varepsilon_{yt} \tag{D.7}$$

   with $Z_{syt} \sim N\left(\mu_{at}, \sigma^2_{at} + \frac{\sigma^2_{z_s}\sigma^2_y}{\sigma^2_{z_s} + \sigma^2_y}\right)$.[4] 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. At the exact time that z_{bt} and Z_{syt} are revealed, the Board and Shareholders may disagree about CEO skill. This timing convention is key for determining SOP outcomes.

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

4. **Board and shareholder beliefs after the annual shareholder meeting and release of 10-K**

   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. A.4.2.

A.4.4 Assumptions about Shareholder strategy

See Section 2.3.3 for the full discussion of the Shareholder's strategy. A primary goal of this paper is to model how the threat of SOP failure influences the Board's wage decision. I model this threat as an *ex ante* probability that the SOP will fail, which is increasing in the wage. In this spirit, the first assumption specifies precisely when shareholders commit to this threat.

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

S must set their probability of failure *before* they receive their private signal or see wages and productivity. This threat of vote failure influences the Board's wage decision. 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.[5]

Assumption 2. *Shareholders seek to optimize operating income, given their beliefs about CEO ability.*

This assumption is a main source of disagreement about CEO ability between the Board and

[5] Based on Assumption 1, Figure A.5 provides a more detailed version of the model timeline (slightly adapting Figure 2.3). 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.

shareholders. The Board's signal z_{bt} causes B to update their beliefs about CEO ability. At the time that the Shareholder commits to their voting strategy, B and S hold different beliefs about CEO ability. This assumption means that S wants to force the Board to pay a wage closer to the unbiased wage, given their beliefs after receiving the signal at the shareholder meeting. Technically, this means that shareholders choose a single s_t for the Board's entire wage schedule — S influence expected wages. Or, put equivalently, S sets an average (across the distribution of z_{bt}) probability of SOP failure that the shareholder base is comfortable with.

Assumption 3. *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 Figure A.4, or Fisch et al., 2018; Novick, 2019, 2020). Further, Table A.2 shows that lagged CEO pay does not influence SOP outcomes. This makes the solution method much simpler, as it avoids higher-order beliefs and an infinite-regress problem in B and S forecasting each other's beliefs (Foster and Viswanathan, 1996; Bonatti et al., 2017).

A.4.5 Full Derivation of Model Solution

Proposition A.4.4. *The Board's problem can be written as*

$$V(\mu_a, \tau, w_{-1}) = \max_{w(s)} \quad \exp\left(\mu_{b|z_b} + 0.5\left(\sigma^2_{b|z_b} + \sigma^2_y\right)\right) w(s)^\alpha - (1-\lambda)w(s) - \chi_B F^U_{\tilde{z}_s}(s \times w(z_b, s)) - AC(w(s), w_{-1}; \tau) + \delta_B\Big[f(\tau_t)V^R + (1 - f(\tau_t))E_{b|z_b}\left[V(\mu'_a, \tau+1, w(s))\right]\Big] \quad \text{(D.8)}$$

where

- $\mu_{b|z_b}$ *and* $\sigma^2_{b|z_b}$ *are defined in* (D.5) *and* (D.6),
- $E_{b|z_b}[V] = F^U_{\tilde{z}_s}(s \times w)$ *where* $F^U_{\tilde{z}_s}$ *is the CDF of the following distribution (see 2.15)*

$$w^U \sim \log N\left(\frac{\mu_s}{1-\alpha} + C, \frac{\sigma^2_s}{(1-\alpha)^2}\right), \quad C = \frac{\log\alpha + \frac{1}{2}\sigma^2_y}{1-\alpha}$$

- $AC(w, w_{-1}; \tau)$ *(adjustment cost) is defined in* (2.3),
- $f(\tau_t)$ *are CEO tenure-specific hazard rates, with* $f_0 = 0$ *and* $f(\tau_t) = 1$
- $V^R = V(\mu_0, 0, 0)$ *as in* (2.20)
- $\mu'_a \mid \mu_{b|z_b}, \tau \sim N\left(\mu_{b|z_b}, \sigma^2_{b|z_b} - \sigma^2_a(\tau + 1)\right)$ *from Prop. A.4.3*

The shareholder's problem can be written as

$$\max_s \int_{z_b} f(z_b \mid \mu_a, \sigma^2_a) \left[\exp\left(\mu_a + 0.5\left(\sigma^2_{b|z_b} + \sigma^2_y\right)\right) w(z_b, s)^\alpha - w(z_b, s) - \chi_S CDF^U_s(s \times w(z_b, s)) \right] dz_b \tag{D.9}$$

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 (D.8). $\mu_{b|z_b}$ 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 2.3.3, but as brief overview it is given by the CDF $F^{\mathrm{U}}_{\tilde{z}_s}$, 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 $N(\mu_0, \sigma^2_0)$ for any state. Hence, it is a boundary condition.

The distribution of μ'_a conditional on μ_a and τ is given in Prop. A.4.3. 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 (D.9), the objects are the same as the Board's problem, however the shareholders choose s under the distribution of z_b, while holding belief of average CEO ability μ_a. 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. Under Assumption 3, shareholders do not behave dynamically, and only vote on the current period. Crucially, as the wage, the productivity signal and the shareholder's private signal are all revealed simultaneously (see Section 2.3.2 and Figures 2.3 and A.5), the Board and shareholders *disagree* about CEO ability at the point the vote is held. In other words, they hold different beliefs about (mean) CEO ability.

The solution $(w(z_b), s)$ is to be found numerically, each $(w(z_b), s)$ is a best response in equilibrium 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 (D.9) for their strategy to be a best response; they cannot deviate at the vote and play a lower threshold.

A.5 Estimation Appendix

A.5.1 Identifying the CEO component of output

To undertake the estimation of CEO skill and output parameters, I need to estimate the CEO component of company output. This process is similar to the analyses undertaken in Matveyev (2017) and Lyman (2023), and the ultimate goal is to give an approximation of how CEO skill (productivity) and effort (via the wage) impact revenues.

I first specify the following functional form for company revenues

$$\log Y_{it} = \underbrace{\log \eta}_{\text{scaling factor}} + \underbrace{a_i + \alpha \log w_{it}}_{\text{CEO component}} + \underbrace{\kappa_1 \log K_{it} + \kappa_2 \log L_{it} + \mu_{IND} + \mu_t}_{\text{non-CEO component}} + \varepsilon_{yit},$$

where i refers to a firm-CEO match, Y_{it} is the revenue generated by the firm, η is a scaling factor, w_{it} is the wage paid to the CEO, a_i is a CEO fixed effect, K_{it} is (PPE), L_{it} is the number of employees, and μ_{IND} and μ_t are industry and time fixed effects. I residualize revenues by netting out the non-CEO component of revenue.

This gives the following form for revenue

$$\log y_{it} = \log \eta + a_i + \alpha \log w_{it} + \varepsilon_{yit} \implies$$

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

The parameter η is a time-invariant scaling factor which scales operating income to its appropriate level. It can also be thought of as how well the firm is able to translate the CEO's skill and effort into actual revenues, and it allows me to translate how CEO effort and wages ultimately affect revenues. While η is not essential for solving the model, it is important in the estimation as it allows for appropriate comparison across subsamples (Page, 2018).

A.5.2 Bunching Estimator Used in the Estimation

This appendix describes the bunching estimator used in the estimation (see Section 2.4.1). It is a simple, reduced-form measure that is simple to implement. Informally, the measure captures how much extra mass is shifted below the failure threshold, relative to a counterfactual density. To start, I bin the data into widths of 0.5% and index each bin by b. I label the bin mid-point as x_b. N_b is the number of SOP vote outcomes that fall in bin b. To estimate a counterfactual density of vote outcomes absent the failure threshold, I first fix a region around the threshold, $\mathbf{E}$, which captures where bunching is most likely to occur

$$\mathbf{E} = \left\{ b \quad \text{s.t.} \quad x_b \in [-e, e] \right\},$$

in the estimation I set $e = 2.5\%$. Excluding $\mathbf{E}$, I estimate the predicted bin counts as a polynomial of the bin midpoints,

$$N_b = \sum_{m=0}^{M} \beta_m (x_b^m)$$

and then predict $\hat{N}_b$ for bins in $\mathbf{E}$. Thus, $\hat{N}_{b|b\in\mathbf{E}}$ represent a counterfactual count for the number of SOP votes that would fall in bin b given there was no bunching in the data. I use these predicted counts to construct the bunching estimator

$$\mathbf{B} = \underbrace{\frac{\sum_{b\in\mathbf{E}} \left(N_b - \hat{N}_b \right) \times \left(\mathbf{1}[x_b < 0] - \mathbf{1}[x_b \geq 0] \right)}{\sum_{b\in\mathbf{E}} \hat{N}_b}}_{\text{Proportion of mass shifted to below threshold}}. \tag{E.1}$$

B measures how many SOP votes are shifted from above the failure threshold to below, expressed as a proportion of the sum of the predicted counts. (E.1) is easily estimable on model-simulated data, so can be used as a targeted moment during the estimation step.

A.5.3 Numerical Solution

The model requires 11 parameters, along with $\mathbf{T} + 1$ 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 10 parameters as Θ,

$$\Theta = \begin{pmatrix} \mu_0 & \sigma_0 & \sigma_y & \alpha & c_W & \sigma_{z_b} & \sigma_{z_s} & \lambda & \chi_B & \chi_S \end{pmatrix}$$

The model's solution proceeds as such

1. Start with a given Θ
2. Discretize each idiosyncratic shock into an N_z grid. E.g., fix the possible realizations of $\varepsilon_{z_{bt}}$, $\varepsilon_{z_{st}}$, etc.
3. Discretize the state space into a $(N_u, \mathbf{T} + 1, N_w)$ grid, call it $\mathbf{S}$, 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 current 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 and Prop. A.4.2 to estimate the continuation value for each tuple (μ, τ, w)
6. For each element in $\mathbf{S}$, solve (D.8) and (D.9),
    - If i % 10 = 0

- Solve B's optimal choice of w' given possible realizations of ε_{z_b}, and how this w' affects the continuation value and adjustment cost.
- Concurrently backing out S' optimal choice of s for the tuple (μ, τ, w) given B's optimal choices of w'
- Update the guess of the value function

- Else,
  - Update B's optimal choices w'
  - Update the guess of the value function

7. Return to 4 and repeat until $\max|V_i - V_{i-1}| < \epsilon = 1e-5$

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

A.5.4 Simulation

I set N_f firms, where N_f is chosen to match the number of firms in the data used for estimation. Given $a \sim (\mu_0, \sigma_0)$, 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}. Given the state, I use the policies described in Section A.5.3 to generate optimal choices. Beliefs update given realizations of σ_{z_b}, σ_{z_s} and z_{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 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.

A.5.5 Estimation

I estimate the 10 parameters

$$\Theta = \begin{pmatrix} \log \eta & \sigma_0 & \sigma_y & \alpha & c_W & \sigma_{z_b} & \sigma_{z_s} & \lambda & \chi_B & \chi_S \end{pmatrix}$$

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 Θ of parameters that minimizes the weighting distance between a vector of moments produced by the model and the corresponding moments computed in the data. That is, given model moment $m(\Theta)$ and data moments $m(X)$ and an appropriate weighting matrix W, I minimize

$$\min_{\Theta} \Big[d(\Theta, X) \Big]' W \Big[d(\Theta, X) \Big] \tag{E.2}$$

$d(\Theta, X)$ is a 13×1 vector of differences between model-simulated and empirical moments (Barrero, 2022). I set $d(\Theta, x) = m(\Theta) - m(X)$. The weighting matrix is the identity matrix. The model moments can be expressed as the coefficients from the following system of equations, where each regression is indexed by i, c, τ, t

$$\begin{aligned}
\ln y &= y_0 + y_1 \ln \text{wage} + e_y \\
(E_{\text{CEO}}[\ln y] - E[E_{\text{CEO}}[\ln y]])^2 &= Var(E[\ln y]) + e_{\text{CEO}} \\
\ln \text{wage} &= b_0 + b_1 \mathbf{1}[\text{SOP fail}] + b_2 \ln \text{lagged wage} + e_b \\
\mathbf{1}[\text{SOP fail}] &= s_0 + s_1 \ln \text{wage} + s_2 \epsilon^y + e_s \\
(\epsilon^y)^2 &= Var(\epsilon^y) + e_{\epsilon_y} \\
\left(\epsilon^b\right)^2 &= Var\left(\epsilon^b\right) + e_{\epsilon_b} \\
(\epsilon^s)^2 &= Var(\epsilon^s) + e_{\epsilon_s} \\
y_{\mathbf{B}} &= \mathbf{B} x_{\mathbf{B}} + e_{\mathbf{B}}
\end{aligned} \tag{E.3}$$

For the final moment **B**, I reconstruct (E.1) using a regression specification.[6] In total, there are 13 moments to pin down 10 parameters. I estimate (E.3) jointly, with standard errors clusted at the CEO-spell level, and use the variance-covariance matrix of the moments to estimate standard errors of the parameter and moments. The full list of moments with notation is displayed in Table A.7.

[6] That is $x_{i,\mathbf{B}} = \sqrt{\frac{\hat{N}_b}{N_b}}$ and $y_{i,\mathbf{B}} = \left(x_{i,\mathbf{B}}\right)^{-1} (\mathbf{1}[x_b < 0] - \mathbf{1}[x_b \geq 0]) \left(1 - \frac{\hat{N}_b}{N_b}\right)$. This expression reconstructs (E.1) in terms of a regression, and conveniently allows me to include (E.1) in our regression system.

A.5.6 Optimization algorithm

My goal is to find the global minimum of the SMM/GMM objective function described in Section A.5.5. 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).[7] 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**

    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 $N = 5000$.

    iii. For each $n \in N$ of the Sobol points, use a minimizer to find the local minimum of each point. Keep the portion p of the points with the smallest local minima to be used in the global stage. In my setup, I use Nelder-Mead locally, and keep the top 4% of points, so I am left with $N_p = 200$ "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.

[7] 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 (2021) for a recent example.

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.

A.5.7 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

$$\begin{aligned}\text{Average firm value} = V_0 &= \sum_{t=1}^{\infty} \delta_B^t E_t[y_t - w_t] \\ &= \sum_{t=1}^{\infty} \delta_B^t E_0[y_0 - w_0] \\ &= \sum_{t=1}^{\infty} \delta_B^t E_0[A_0 \times (w_0)^{\alpha} - w_0] \\ &= \frac{1}{1-\delta_B}\left[\exp\left(\mu_0 + 0.5\left(\sigma_0^2 + \sigma_y^2\right)\right) \times (w_0)^{\alpha} - w_0\right]\end{aligned}$$

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

$$\text{SOP failure cost (\% average value)} = \frac{\chi_{\{B,S\}}}{V_0} \tag{E.4}$$

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 $\nu \in [0, 1]$ on the CEO's utility (pure dollar wage), and $1 - \nu$ on company profits, so that their program is

$$\max_{w_t} \quad (1-\nu) \times E_0([A_0]w_t^{\alpha} - w_t) + \nu \times w_t$$

As can be seen, this is the same as $\max_{w_t} E_0[A_0]w_t^{\alpha} - (1-\lambda)w_t$, with $\lambda = \frac{\nu}{1-\nu}$, the main program from the paper. Equivalently defining $\nu = \frac{\lambda}{1+\lambda}$, if $\lambda = 1$, then the Board sets the wage such that company profits go to zero (perfect capture). If $\lambda = 0$, the Board maximizes profits. The optimal split that the Board decides for the CEO is thus

$$\theta_{\text{CEO}} = \frac{\nu w_t}{(1-\nu) \times (E_0[A_0]w_t^{\alpha} - w_t) + \nu w_t} = \frac{\lambda w_t}{E_0[A_0]w_t^{\alpha} - (1-\lambda)w_t}$$

The split θ_{CEO} describes how much the Board tilts the surplus towards the CEO. I can thus describe that the surplus split for the average CEO as

$$\theta_{\text{CEO}} = E_0\left[\frac{\lambda w_t}{y_t - (1-\lambda)w_t}\right] = \frac{\lambda w_0}{y_0 - (1-\lambda)w_0} \tag{E.5}$$

A.5.8 The Impact of the Shareholder SOP Failure Cost

In companion analysis to Section 2.5.2, I evaluate the effect of removing the shareholder cost to SOP failure. Removing this shareholder cost represents a *significant* change to the underlying preferences of shareholders. The exercise of changing χ_S while holding other parameters constant is likely not valid. By its nature, SOP is akin to a costly punishment mechanism. In some sense, the cost to shareholders has emerged endogenously as a way validate SOP, even though it is a voice mechanism. By failing the SOP, shareholder *commit* to re-considering the careers of directors (and possibly the CEO); these outcomes incur costs. If SOP were truly costless to shareholders in reality, then the Board would likely view it as "cheap talk." Nonetheless, understanding its impact on CEO pay and firm value is insightful for the model's predictions.

Table A.8 displays the results. When $\chi_S = 0$, the SOP failure rate increases from about 7% to 42%. The reason the failure rate does not go beyond this can be gleaned from Figure 2.4; shareholders do not want to decrease below the unbiased wage, and for each state there is a uniquely-defined probability of failure that maximizes Shareholder welfare.

CEO wages would fall by 4.7% on average, and at the 25th percentile wages would decrease by 8.6%. Firm value value increases, but marginally. By enforcing a much higher degree of SOP failure, there are states when shareholders over-discipline the Board. That is, while wages being kept low for bad CEOs is good, it may often be the case that the Board's signal means the CEO is talented and should be paid more. It is worth commenting on the relative value effects in Table 2.7 (removing the Board cost to SOP failure) and Table A.8 (removing the Shareholder cost). The estimation shows that the Board cost is larger: equivalent to 2.06% of value, relative to 0.76% for shareholders. The relative value effect of these two counterfactuals depends on the relative magnitude of these cost parameters.

Summing up, this counterfactual shows that, while the shareholder cost to SOP failure keeps the SOP failure rate low and leads to the Board paying the CEO a higher wage, the impact of

removing this cost on value is moderate. SOP is already effective at disciplining the Board — SOP failure is costly enough to the Board to keep wages down. My analysis in Section 2.6 shows that changing the structure of SOP to give shareholders more input into the compensation decision would lower wages (on average) and positively impact firm value. Further, removing this cost is likely too drastic of a change to preferences, and probably renders this counterfactual invalid.

Appendix B

Appendix to Chapter 3

B.1 Appendix Figures and Tables

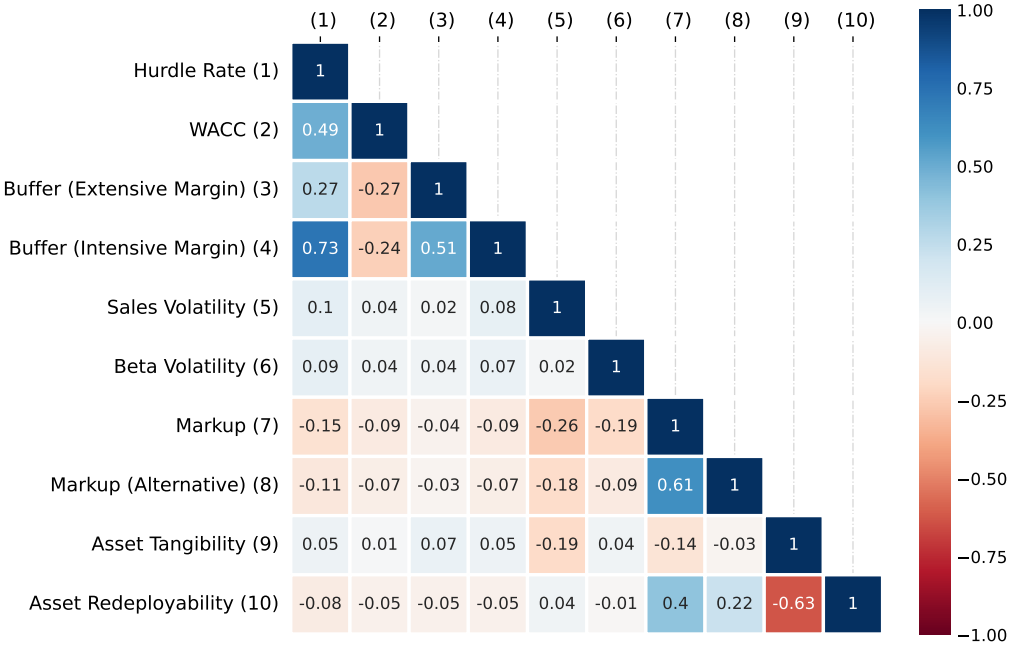

FIGURE B.1: Cross-Correlations of Main Variables

This figure shows the pair-wise correlations among the main variables. Dark blue (red) indicates strong positive (negative) correlations. Detailed variable definitions are available in Table B.1.

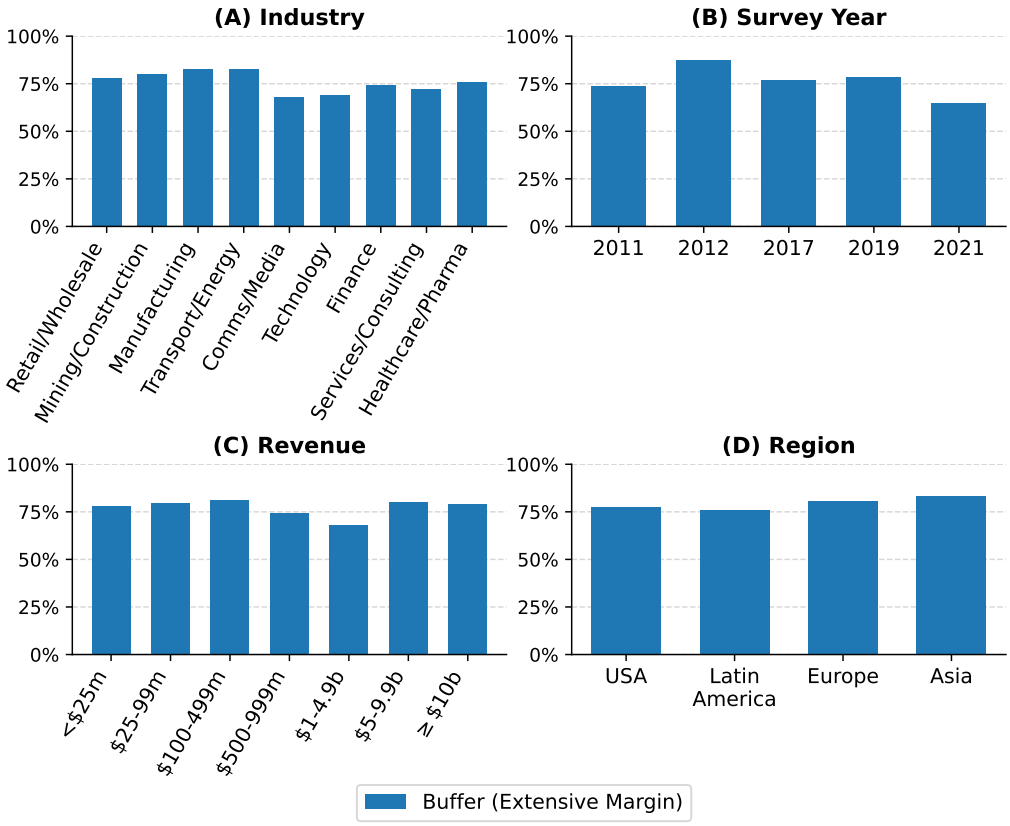

FIGURE B.2: Time Series and Cross-Sectional Differences in the Extensive Margin of the Buffer

This figure displays the proportion of firms that use a positive buffer (i.e., their hurdle rate exceeds their cost of capital) across several observable characteristics of the CFO survey sample.

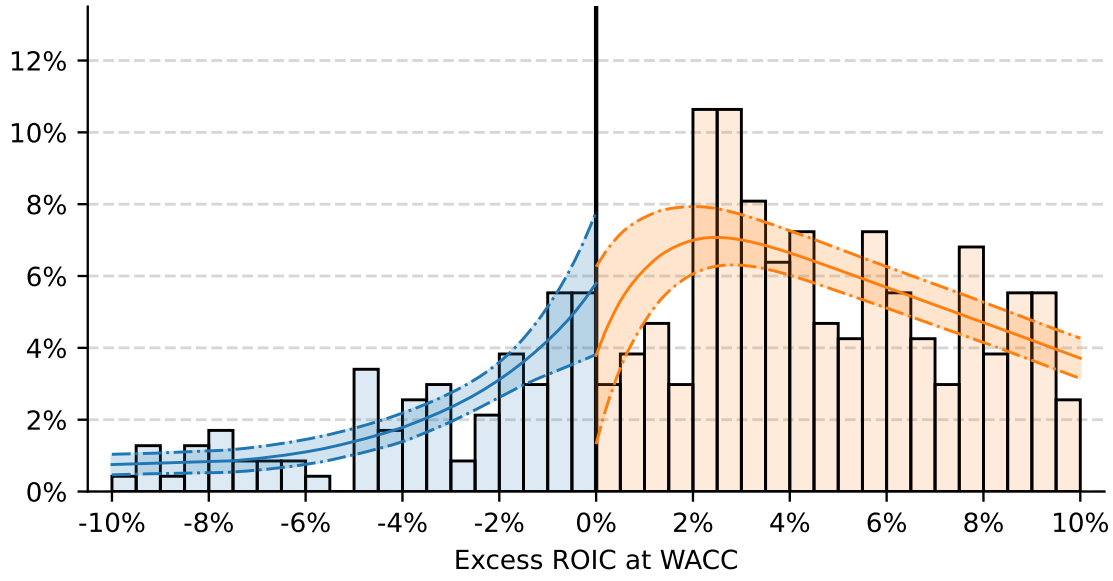

FIGURE B.3: Density Manipulation in Excess ROIC Relative to WACC

This figure displays robustness for Figure 3.4. In particular, we test for density manipulation of ROIC around the cost of capital, as opposed to the hurdle rate. The sample is the same as in Figure 3.4. We define $Excess\ ROIC_{i,t}^{WACC} = ROIC_{i,t} - WACC_{i,t}$ as excess ROIC relative to the firm's cost of capital (as opposed to the hurdle rate as in Figure 3.4). We implement the same test and look for excess bunching in ROIC directly above the firm's WACC; we do not find a discontinuity in the density at WACC (the test returns a t-statistic of −1.013).

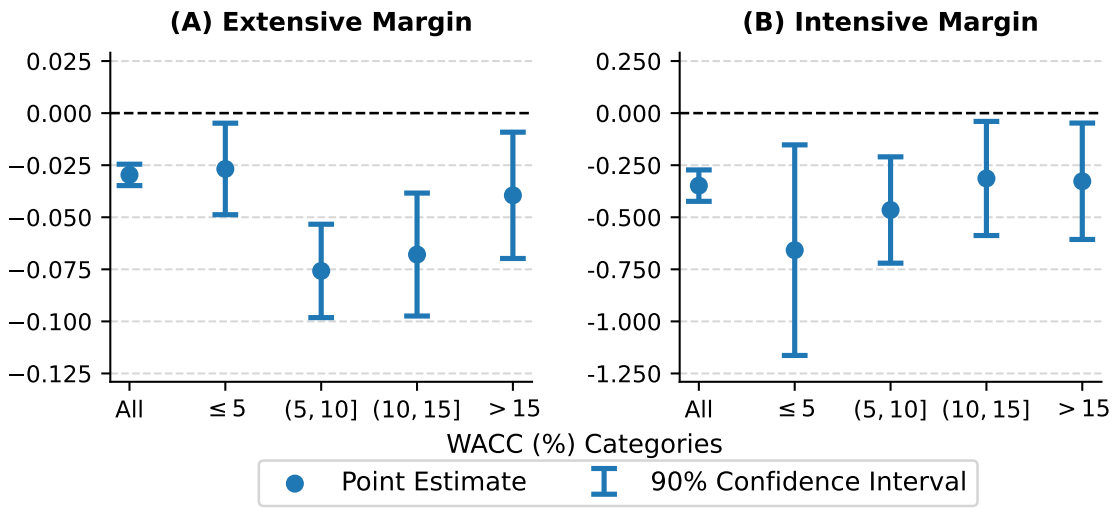

FIGURE B.4: Negative Relation Between Buffer and Cost of Capital Across Different WACC/Risk Groups

This figure displays how the relation between the cost of capital and the buffer varies across different levels of the cost of capital. We first estimate a regression of the buffer (extensive and intensive margins) on an interaction between WACC and a categorical variable that groups survey firms by the level of their WACC (along with the non-interactive variables):

$$B_{i,t} = \delta_1 + \beta_1 WACC_{i,t} + \sum_{k=2}^{4} \beta_k WACC_{i,t} \times WACC\ \text{Group}_k + \delta_k WACC\ \text{Group}_k + X_{i,t}\gamma'_{i,t} + \varepsilon_{i,t}$$

The four groups are (1) ≤ 5%, (2) (5%, 10%], (3) (10%, 15%] and (4) > 15%. We then display effect of WACC on the buffer for each group. For example, for firms with WACC ≤ 5%, the figure displays the baseline effect from the regression (β_1); the other three groups display the group-specific slope coefficient ($\beta_1 + \beta_{k|k\in\{2,3,4\}}$). In each plot, we display the overall effect for comparison (labeled "All"), which is estimated from a similar regression without interactions. Controls in this regression are sales volatility and beta volatility; fixed effects are survey-quarter and size. Standard errors are clustered by survey industry × survey quarter. The results reveal a significant negative relation between the hurdle rate buffer and WACC for each subgroup.

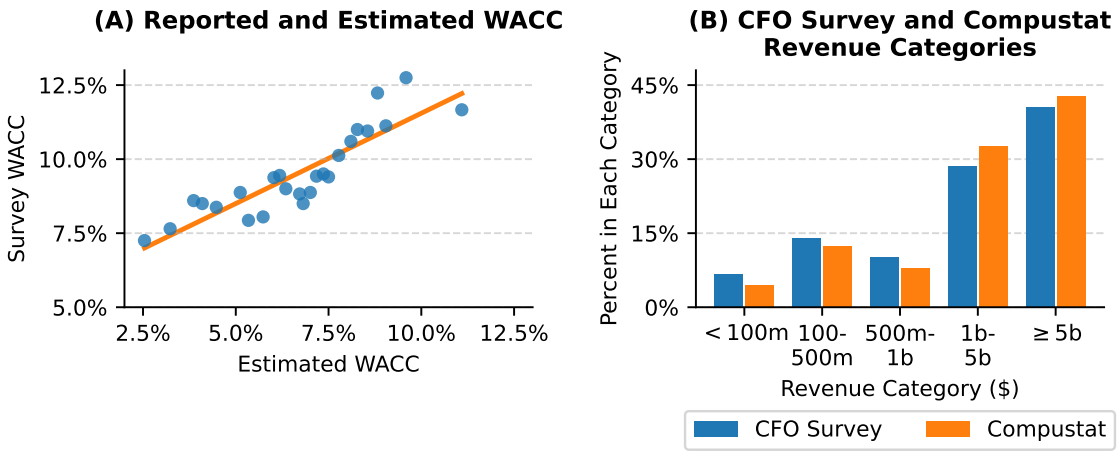

FIGURE B.5: Closeness of Data Reported on CFO Survey to Archival Data

This figure analyzes the closeness of survey-reported data and archival data using CFO survey firms that appear in Compustat. We match data from Compustat to the survey firms by survey year and calendar year in Comustat. Panel A displays a binned scatter plot from a regression of reported WACC on estimated WACC. The cost of equity is estimated using the CAPM, cost of debt is total interest expense to total debt. Panel B displays the percentage of these Compustat firms that fall into the stated revenue category on the survey (blue bars) *vs.* the same categories calculated using archival revenue data (orange bars).

Table B.1: Empirical Variable Definitions

Variable	Description	Source
Hurdle Rate	The minimum rate of return required to pursue a project	CFO Survey
Cost of Capital/WACC	The firm's weighted average cost of capital	CFO Survey
Buffer (Extensive Margin)	Indicator variable taking a value of one if the firm's hurdle rate exceeds its cost of capital; i.e., the extensive margin of the buffer	CFO Survey
Buffer (Intensive Margin)	The difference between hurdle and WACC, i.e. the intensive margin of the buffer	CFO Survey
Buffer | Buffer > 0	The intensive margin of the buffer for firms that use a positive buffer. This differs from the definition above as it excludes Buffer = 0 observations	CFO Survey
Survey Industry	Industry of firm, as supplied by respondent to the CFO survey, roughly equivalent to 1-digit SIC. The 10 categories are Retail/Wholesale Trade, Mining/Construction, Manufacturing, Transportation/Energy, Communications/Media, Technology, Finance, Servces/Consulting, Healthcare/Pharma, Other.	CFO Survey
Has Credit Rating	Indicator variable equal to one if the firm states it has a credit rating (and zero for all other firms)	CFO Survey
Size	Categorical variable tracking firm size by revenue. Categories are $< \$25m$, $\$25 - 99m$, $\$100 - 499m$, $\$500 - 900m$, $\$1 - 4.9b$, $\$5 - 9.9b$ and $\geq \$10b$. Often included as a fixed effect in analysis	CFO Survey
Large Firm	Indicator variable equal to one if the firm's sales revenue is weakly greater than $1 billion	CFO Survey
Public Firm	Indicator variable equal to one if the firm is publicly traded	CFO Survey
Firm-Level (US) Optimism	Answer to the question: "Rate your optimism about your firm (or in a separate question, the US economy) on a scale from 0-100, with 0 being the least optimistic and 100 being the most optimistic."	CFO Survey
Return on Invested Capital	ROIC is $EBIT_t \times (1 - tax_t)/ICAPT_{t-1}$ where $ICAPT$ is Invested Capital. Invested Capital is the sum of long-term debt (total) and common equity (total). We take trailing four-quarter sums of $EBIT(1 - tax)$, and divide by the invested capital from the end of the previous fiscal year to derive our final measure.	Compustat
Beta Volatility	Standard deviation of the CAPM beta within a quarter and 4-digit NAICS industry	Compustat & CRSP

Sales Volatility	We adapt methodology in Décaire (2021) to estimate industry idiosyncratic sales volatility. We estimate $sales/(lagged\ assets)_{i,j,t,q} = \alpha_j + \alpha_t + \alpha_q + \varepsilon_{i,j,t,q}$ for firm i, industry j, calendar quarter t and $q \in \{1,2,3,4\}$. We take residuals $\varepsilon_{i,j,t,q}$ as a quarterly firm-level measure of idiosyncratic risk. To net out short-run, firm-specific shocks, we take trailing 2-year averages of firm idiosyncratic risk as our final firm-level measure. The final measure is the standard deviation of the firm-level measure for firms in a given NAICS-4 industry and calendar quarter.	Compustat
Customer Concentration	Average of firms' corporate customer sales HHI within a year and NAICS-3 industry. For firm i and customer firm(s) j, $HHI_i^{sales} = \sum_j \left(\frac{sales_{i,j}}{sales_i}\right)^2$. To measure HHI, we use the Compustat Segments Customers file (which itself is created using SEC filings), focusing on corporate customers. See, e.g., Patatoukas (2012), for more details.	Compustat & SEC filings
Markup	We follow the "accounting markup" from Baqaee and Farhi (2020b) (BF). At the firm-level, we measure profits as operating income after depreciation (OIBPD-DP). The firm-level markup is μ_i comes from the following relation: $\text{profits}_i = \left(1 - \mu_i^{-1}\right)\text{sales}_i$. To net out the effects of year-specific shocks, we take four-year trailing averages of μ_i within-firm. Following BF, we take NAICS-3 by year averages of μ_i as our industry measure. See Appendix C.2.1 of BF details, and Baqaee and Farhi (2019) for details on replication.	Compustat
Markup (Alternative)	We follow the estimation procedure PF1 from De Loecker et al. (2020), with the alterations described in Appendix C.2.3 of BF. We estimate output elasticities at the NAICS-3 level using Compustat data from 1990-2021, then calculate firm-level markups and take trailing four-year averages to net out year-specific shocks. Finally, we take NAICS-3 by year averages of markups as our industry measure. See Baqaee and Farhi (2019) for details on replication.	Compustat & FRED
Consumer-Facing	Indicator variable taking a value of one if a firm is in GICS sectors 25 or 30 ("Consumer Discretionary" or "Consumer Staples")	GICS
Asset Tangibility	We calculate firm-level tangibility as the ratio of property, plant and equipment to total fixed assets (PPENT/AT). Because there is very little time variation in tangibility (a regression of firm tangibility on NAICS-4 fixed effects produces an R^2 of 0.62, adding year fixed effects produces an R^2 of 0.635), we take simple NAICS-4 averages of firm-level tangibility for our industry measure.	Compustat
Asset Redeployability	See Kim and Kung (2017). Asset redeployability (at the asset level) is the proportion of industries that use a given asset. Industry-level asset redeployability is the value-weighted average of redeployability based on how important assets are to an industry. We use the publicly available data from Kim and Kung (2017) and match to our survey firms at the NAICS-2 level.	Kim and Kung (2017), BEA

This table gives definitions, details on construction and sources for each empirical variable used in the paper. Appendix B.3 contains additional details on the survey data and other data sources.

Table B.2: Reason Aggregation for Figure 3.1

Panel A: 2011

Financing Constraint	1.	Shortage of funding
Managerial/Resource Constraint	1.	Shortage of employees
	2.	Shortage of management time and expertise
	3.	Shortage of production capacity
Project Prioritization	1.	We do not pursue some positive net present value projects because we think others will earn even higher returns
Over-Optimism/Agency	1.	Some projects only appear to be attractive due to optimistic projections but may not be successful
	2.	Project might reduce earnings per share
Idiosyncratic Risk/Uncertainty	1.	There is too much uncertainty about some projects
	2.	The risk of the project is too high

Panel B: 2019

Financing Constraint	1.	Because we face funding constraints
Managerial/Resource Constraint	1.	Because we have scarcity of managerial time/expertise
Project Prioritization	1.	So that we choose only the best available projects
Over-Optimism/Agency	1.	So that we choose projects that are profitable
	2.	To provide a buffer in case the project underperforms
Idiosyncratic Risk/Uncertainty	1.	To account for riskiness of the projects being evaluated
	2.	To account for costs not captured by WACC
	3.	To provide a margin of error in calculations and assumptions

Panel C: 2022

Financing Constraint	1. Our firm cannot fund all profitable projects
Managerial/Resource Constraint	1. Scarcity of non-management labor 2. Scarcity of management times
Project Prioritization	1. To limit the total number of projects we take on 2. Saves resources in order to preserve the option to invest in future projects that might earn higher return
Over-Optimism/Agency	1. Helps offset possible over-optimism in project evaluation 2. Builds in a buffer, to reduce the odds we will have to cancel a project that we have already started should a negative surprise occur
Idiosyncratic Risk/Uncertainty	1. Provides a margin of error in calculations and assumptions 2. Accounts for project-specific risks not reflected in WACC

This table displays how we produce Figure 3.1. The possible responses that CFOs can give for the reason(s) that they set their hurdle above WACC vary by survey. We categorize possible reasons into five qualitative groupings.

Table B.3: Hurdle Rates, Costs of Capital and Buffers by Industry

Panel A: NAICS-2 Averages

Code	Description	Percent of Sample	Hurdle	WACC	Buffer		
					Intensive Margin	Extensive Margin	Buffer | Buffer > 0
11	Agri, Forestry, Fish	0.32	11.875	9.250	2.625	0.750	3.500
21	Mining, Oil/Gas	2.03	18.020	10.140	7.880	0.960	8.208
22	Utilities	0.97	8.892	7.075	1.817	0.667	2.725
23	Construction	5.19	13.575	8.322	5.253	0.734	7.153
31-33	Manufacturing	21.59	14.616	9.197	5.418	0.805	6.735
41-42	Wholesale Trade	7.22	13.553	8.576	4.977	0.775	6.419
44-45	Retail Trade	5.11	14.374	8.910	5.465	0.778	7.026
48-49	Transport/Warehouse	1.95	13.562	9.017	4.546	0.833	5.455
51	Information	4.38	15.398	9.645	5.753	0.778	7.397
52	Finance/Insurance	9.42	11.693	7.723	3.970	0.733	5.418
53	Real Estate	3.9	12.964	8.383	4.580	0.729	6.281
54	Professional/Science Svcs.	9.66	15.708	9.936	5.771	0.756	7.631
55	Mgmt. of Companies	0.73	10.267	6.839	3.428	0.556	6.170
56	Admin/Waste Mgmt.	2.76	15.294	8.751	6.543	0.824	7.945
61	Educational Svcs.	1.38	8.444	6.368	2.076	0.529	3.922
62	Healthcare	3.98	11.849	7.658	4.191	0.837	5.009
71	Arts/Entertainment	0.49	11.383	9.550	1.833	0.667	2.750
72	Accommodation/Food	1.46	16.744	9.386	7.358	0.833	8.830
81,91-92	Other	2.68	11.955	6.564	5.391	0.848	6.354
Unknown	Unknown	14.77	13.848	8.903	4.945	0.775	6.383

Panel B: Survey Industry Averages

Industry	Percent of Sample	Hurdle	WACC	Buffer		
				Intensive Margin	Extensive Margin	Buffer | Buffer > 0
Mining/Construction	4.87	14.212	7.894	6.317	0.800	7.897
Technology	7.31	16.394	10.638	5.756	0.689	8.356
Manufacturing	24.11	14.803	9.276	5.527	0.828	6.673
Communications/Media	2.52	15.581	10.165	5.416	0.677	7.995
Services/Consulting	12.01	13.764	8.664	5.100	0.723	7.055
Retail/Wholesale	10.39	13.170	8.320	4.851	0.781	6.209
Transportation/Energy	7.87	13.349	8.654	4.696	0.825	5.694
Healthcare/Pharma	6.09	13.128	8.763	4.365	0.760	5.744
Finance	12.42	11.958	7.790	4.168	0.745	5.594
Other	12.42	13.459	8.274	5.185	0.797	6.503

This table displays industry-level averages of hurdle rates, costs of capital and the intensive and extensive margins of the buffer. Panel A sorts by NAICS-2 industries, Panel B sorts by survey industries (i.e., the industry the CFO stated on the survey), which approximately align with one-digit SIC codes. The final column of both panels displays the average buffer conditional on using a positive buffer. See Table B.1 for detailed definitions.

Table B.4: Excess ROIC Density Manipulation Test Robustness

		(1)	(2)	(3)	(4)	(5)
			Bandwidth			
		Data-Driven	2.5	5	7.5	10
Polynomial Order	1	0.057^{***}	0.083^{***}	0.030^{*}	0.004	0.002
		(0.021)	(0.027)	(0.017)	(0.012)	(0.010)
		$\{4.422, 3.432\}$				
	2	0.051^{**}	0.108^{**}	0.097^{***}	0.036^{*}	0.012
		(0.023)	(0.048)	(0.028)	(0.021)	(0.017)
		$\{7.926, 6.038\}$				
	3	0.125^{***}	0.136^{*}	0.118^{***}	0.113^{***}	0.046^{*}
		(0.034)	(0.077)	(0.041)	(0.030)	(0.025)
		$\{7.449, 6.199\}$				

This table displays robustness the results of the density manipulation test described in Figure 3.4. We alter both order of the local polynomial estimator, as well as the bandwidth for the local estimation. For each polynomial order p, the bias-corrected density estimator is set to be $q = p + 1$, which is the default setting. In each column, we display the difference in the estimated density immediately above and immediately below the cutoff. In the notation of Cattaneo et al. (2018), we display $\hat{f}_{+,p}(h) - \hat{f}_{-,p}(h)$, where $\hat{f}_{\pm,p}(h)$ is derived via local polynomial density estimation. Below the estimated difference in density, we display standard errors. Column 1 displays the results when the bandwidths for estimation are chosen via the data-driven methodology in Section 2.5 of Cattaneo et al. (2018). In column 1, the chosen bandwidths $\{h_-, h_+\}$ are displayed below the standard error. In columns 2-5, we fix the bandwidth to be equal on both sides. ***, ** , * denote significance at 1%, 5%, 10%.

Table B.5: Additional Analysis of the Relation Between The Buffer and Price Markups

Panel A: The Buffer and Markups for All CFO Survey Firms

	(1)	(2)	(3)	(4)	(5)	(6)
	Extensive Margin			Intensive Margin		
Markup	-0.025**	-0.019	-0.021	-0.623***	-0.466***	-0.505***
	(0.011)	(0.012)	(0.014)	(0.139)	(0.154)	(0.164)
Cost of Capital	-0.028***	-0.028***	-0.029***	-0.320***	-0.324***	-0.322***
	(0.003)	(0.003)	(0.003)	(0.045)	(0.045)	(0.045)
Sales Volatility		0.003	0.003		0.303**	0.236
		(0.015)	(0.014)		(0.151)	(0.162)
Beta Volatility		0.027**	0.027**		0.378**	0.336**
		(0.012)	(0.012)		(0.149)	(0.144)
Observations	1,232	1,232	1,232	1,232	1,232	1,232
R-squared	0.078	0.082	0.124	0.071	0.079	0.107
Survey Quarter FE			Yes			Yes
Size FE			Yes			Yes
Ownership FE			Yes			Yes
Credit Rating FE			Yes			Yes

Panel B: The Buffer and Markups for Non-Consumer- and Consumer-Facing Industries

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Extensive Margin				Intensive Margin			
Slope Coefficient								
Firms in Non-Consumer-Facing Industries	-0.031***	-0.024**	-0.027*	-0.026*	-0.696***	-0.526***	-0.569***	-0.559***
	(0.012)	(0.012)	(0.014)	(0.014)	(0.148)	(0.162)	(0.173)	(0.164)
Firms in Consumer-Facing Industries	0.056	0.061	0.052	0.040	0.484	0.590	0.577	0.526
	(0.037)	(0.038)	(0.037)	(0.038)	(0.409)	(0.431)	(0.428)	(0.479)
Observations	1,232	1,232	1,232	1,232	1,232	1,232	1,232	1,232
Controls		Yes	Yes	Yes		Yes	Yes	Yes
Fixed Effects			Yes	Yes			Yes	Yes
Fully Interacted				Yes				Yes

Panel C: Alternative Measure of Markups

	(1)	(2)	(3)	(4)	(5)	(6)
	Extensive Margin			Intensive Margin		
Markup (Alternative)	-0.029**	-0.024**	-0.018	-0.556***	-0.440***	-0.420***
	(0.011)	(0.011)	(0.012)	(0.158)	(0.164)	(0.163)
Cost of Capital	-0.027***	-0.027***	-0.028***	-0.321***	-0.326***	-0.320***
	(0.003)	(0.003)	(0.003)	(0.046)	(0.047)	(0.037)
Beta Volatility		0.029**	0.031**		0.484***	0.476***
		(0.012)	(0.013)		(0.170)	(0.165)
Sales Volatility		0.007	0.009		0.368**	0.316**
		(0.016)	(0.012)		(0.145)	(0.159)
Observations	1,066	1,066	1,066	1,066	1,066	1,066
R-squared	0.082	0.087	0.118	0.073	0.085	0.109
Survey Quarter FE			Yes			Yes
Size FE			Yes			Yes
Ownership FE			Yes			Yes
Credit Rating FE			Yes			Yes

This table provides additional analysis of the relation between hurdle rate buffers and price markups. In Panel A, we repeat the analysis of Table 3.6 for the full sample, i.e., we group non-consumer-facing and consumer-facing together and estimate the relationship between buffers and price markups. Like Table 3.6, the markup is the "accounting" markup from Baqaee and Farhi (2020b). In Panel B, in order to separate out predominantly B2B and predominantly consumer-facing firms, we display the estimated slope coefficient on the markup for firms in non-consumer-facing and consumer-facing industries separately. In columns 2 and 6, we include the controls Beta Volatility and Sales Volatility. In columns 3 and 7, we also include survey quarter, size, ownership and credit rating fixed effects. In columns 4 and 8, we interact all right-hand side variables with the consumer-facing dummy. Lastly, Panel C repeats the analysis of Table 3.6 using an alternative measure of Industry Markups. We estimate output elasticities via the production function estimation approach described in De Loecker et al. (2020), Baqaee and Farhi (2020b) and Olley and Pakes (1996). In Panels A and C, all variables apart from the cost of capital are standardized to unit variance. All variables are defined in detail in Table B.1. In each panel, standard errors are clustered at the survey industry × survey quarter and displayed in parentheses below the coefficient. ***, ** , * denote significance at 1%, 5%, 10%.

Table B.6: The Buffer and Firm-Level Bargaining Power Measures

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Customer Concentration				Markup			
Buffer Margin:	Extensive		Intensive		Extensive		Intensive	
Customer Concentration (Firm-Level)	0.055*	0.050	0.187	0.250				
	(0.033)	(0.034)	(0.391)	(0.394)				
Markup (Firm-Level)					-0.057**	-0.034	-0.859**	-0.742**
					(0.028)	(0.042)	(0.378)	(0.292)
Cost of Capital	-0.025**	-0.028***	-0.351	-0.437**	-0.028***	-0.029***	-0.411*	-0.480**
	(0.011)	(0.009)	(0.212)	(0.216)	(0.010)	(0.009)	(0.218)	(0.220)
Beta Volatility		0.017		0.531		0.020		0.451
		(0.022)		(0.331)		(0.020)		(0.316)
Sales Volatility		0.036		0.484*		0.027		0.254
		(0.030)		(0.265)		(0.031)		(0.255)
Observations	178	178	178	178	178	178	178	178
R-squared	0.040	0.141	0.044	0.136	0.040	0.133	0.074	0.151
Survey Quarter FE		Yes		Yes		Yes		Yes
Size FE		Yes		Yes		Yes		Yes
Ownership FE		Yes		Yes		Yes		Yes
Credit Rating FE		Yes		Yes		Yes		Yes

This table provides further analysis on how bargaining power affects the buffer. We corroborate our findings in Tables 3.5 and 3.6 by focusing on the CFO Survey firms that appear in Compustat, for which we can observe firm-level customer concentration and price markups. We focus on CFO survey firms that are also present in Compustat. For consistency across the specifications in the table, we further focus on firms in non-consumer-facing industries (see Table 3.6). In columns 1-4, we focus on the degree of concentration in a firm's corporate customers as our measure of bargaining power (see Table 3.5). Columns 1-2 focus on the extensive margin of the buffer and columns 3-4 focus on the intensive margin. Similarly, columns 5-8 focus on the price markup for predominantly B2B firms (see Table 3.6). All variables apart from Cost of Capital are standardized to unit variance. Standard errors are clustered at survey industry × survey quarter and displayed in parentheses below the coefficient. ***, ** , * denote significance at 1%, 5%, 10%.

Table B.7: The Buffer and Asset Redeployability (Kim and Kung, 2017)

	(1)	(2)	(3)	(4)	(5)	(6)
	Extensive Margin			Intensive Margin		
Asset Redeployability	-0.026*	-0.026	-0.030*	-0.319**	-0.323**	-0.421**
	(0.016)	(0.016)	(0.015)	(0.148)	(0.143)	(0.159)
Cost of Capital	-0.027***	-0.028***	-0.029***	-0.311***	-0.321***	-0.317***
	(0.003)	(0.003)	(0.003)	(0.044)	(0.045)	(0.044)
Beta Volatility		0.030***	0.030**		0.459***	0.420***
		(0.011)	(0.011)		(0.148)	(0.142)
Sales Volatility		0.008	0.009		0.436***	0.378**
		(0.015)	(0.014)		(0.150)	(0.159)
Observations	1,232	1,232	1,232	1,232	1,232	1,232
R-squared	0.078	0.084	0.127	0.062	0.076	0.105
Survey Quarter FE			Yes			Yes
Size FE			Yes			Yes
Ownership FE			Yes			Yes
Credit Rating FE			Yes			Yes

This table explores how the redeployability of a firm's assets affects the use and size of the buffer. Prices of assets with low redeployability are more likely to be bargained over, hence we predict a negative relationship between the buffer and redeployability, all else equal. Our measure of asset redeployability is the industry-level (BEA) measure from Kim and Kung (2017). We use the BEA-NAICS crosswalk provided with their publicly available data to match the measure to our survey firms at the NAICS-2 level. The mean (standard deviation) asset redeployability for industries in the CFO survey is 0.44 (0.13), which closely matches the broader BEA-level data in Kim and Kung (2017): 0.41 (0.12). Asset Redeployability, Beta Volatility and Sales Volatility are standardized to mean zero, unit variance. Standard errors are clustered at the survey industry × survey quarter and displayed in parentheses below the coefficient. ***, ** , * denote significance at 1%, 5%, 10%.

B.2 Proofs

Proof of Lemmas 1 and 2

Proof 1. If $s = \{\theta(x_F + x_O)R - x_F H_F, (1-\theta)(x_F + x_O)R - x_O H_O\}$, $d = \{0, 0\}$, and $\alpha = \beta = 1$, then (3.1) becomes

$$\max_{\theta} \quad \Big(\theta(x_F + x_O)R - x_F H_F\Big)\Big((1-\theta)(x_F + x_O)R - x_O H_O\Big). \tag{B.1}$$

Expansion yields

$$\theta(1-\theta)[(x_F + x_O)R]^2 - \theta(x_F + x_O)Rx_O H_O - (1-\theta)(x_F + x_O)Rx_F H_F + x_F H_F x_O H_O.$$

Taking first-order conditions,

$$(1-2\theta)[(x_F + x_O)R]^2 + (x_F + x_O)Rx_F H_F - (x_F + x_O)Rx_O H_O = 0.$$

Solving for the Nash split yields

$$\theta = \frac{1}{2} + \frac{x_F H_F - x_O H_O}{2(x_F + x_O)R}.$$

2. If $s = \{\theta[(x_F + x_O)R - (x_F H_F + x_O H_O)], (1-\theta)[(x_F + x_O)R - (x_F H_F + x_O H_O)]\}$, $d = \{0, 0\}$, $\alpha = x_F H_F$, and $\beta = x_O H_O$, then (3.1) becomes

$$\max_{\theta} \quad \Big(\theta[(x_F + x_O)R - (x_F H_F + x_O H_O)]\Big)^{x_F H_F}\Big((1-\theta)[(x_F + x_O)R - (x_F H_F + x_O H_O)]\Big)^{x_O H_O}. \tag{B.2}$$

Taking first-order conditions,

$$\frac{1}{(\theta-1)\theta}\Big[(\theta((R - H_F)x_F + (R - H_O)x_O))^{H_F x_F}$$
$$((\theta-1)((H_F - R)x_F + (H_O - R)x_O))^{H_O x_O}(H_F(\theta-1)x_F + H_O\theta x_O)\Big] = 0 \tag{B.3}$$

Solving for the Nash split yields

$$\theta = \frac{x_F H_F}{x_F H_F + x_O H_O}.$$

Proof of Proposition 1

Proof Integrating (3.6) over all incentive compatible gross returns ($\underline{V}'$ to $\overline{V}$) yields

$$V_B = \frac{x_F \tau}{x_F \tau + x_O W_O} V \Big|_{\underline{V}'}^{\overline{V}}.$$

Substituting in $\underline{V}' = x_F \tau + x_O W_O$ yields the following:

$$V_B = \frac{x_F \tau}{x_F \tau + x_O W_O} [\overline{V} - (x_F \tau + x_O W_O)].$$

Simplifying gives the expression for the value of the firm in terms of τ:

$$V_B = \frac{x_F \tau \overline{V}}{x_F \tau + x_O W_O} - x_F \tau. \tag{B.4}$$

Taking the derivative of (B.4) with respect to τ yields the optimal hurdle rate

$$\begin{aligned} \frac{\partial V_B}{\partial \tau} &= \frac{x_F \overline{V}(x_F \tau + x_O W_O) - x_F^2 \tau \overline{V}}{(x_F \tau + x_O W_O)^2} - x_F = 0 \\ \implies & \overline{V} x_O W_O = (x_F \tau + x_O W_O)^2 \\ \implies & x_F \tau = \sqrt{\overline{V} x_O W_O} - x_O W_O. \end{aligned} \tag{B.5}$$

The optimal τ is therefore given by

$$\tau^* = \frac{\sqrt{\overline{V} x_O W_O} - x_O W_O}{x_F},$$

which is (3.7) in the text. It follows that $\tau^* > W_F$ if (3.10) holds. Otherwise, $\tau^* = W_F$.

Plugging (3.7) into (B.4) for τ gives us the value of the project at the optimal IRR hurdle rate

$$V_B^* = \frac{\overline{V}(\sqrt{\overline{V} x_o W_O} - x_O W_O)}{\sqrt{\overline{V} x_O W_O}} - \sqrt{\overline{V} x_O W_O} + x_O W_O$$

which simplifies to

$$V_B^* = \overline{V} + x_O W_O - 2\sqrt{\overline{V} x_O W_O}\,.$$

This is (3.8) in the text. The change in value, ΔV, from using a buffered hurdle rate is calculated as the difference between V_N from (3.5) and V_B from (3.8)

$$\Delta V = \overline{V} + x_O W_O - 2\sqrt{\overline{V} x_O W_O} - \frac{x_F W_F \overline{V}}{x_F W_F + x_O W_O} + x_F W_F$$

$$\implies = (1 - \frac{x_F W_F}{x_F W_F + x_O W_O})\overline{V} + x_F W_F + x_O W_O - 2\sqrt{\overline{V} x_O W_O}. \tag{B.6}$$

It is straightforward to show that $\Delta V > 0$. We can re-write (B.6) as

$$\Delta V = \frac{x_O W_O \overline{V}}{\underline{V}} + \underline{V} - 2\underline{V}',$$

or

$$\Delta V = (\underline{V}')^2 + \underline{V}^2 - 2\underline{V}'\underline{V}. \tag{B.7}$$

This expression on the right side is of the form $a^2 + b^2 - 2ab$, which is strictly greater than zero for all values of a and b. Therefore, B.7 is always positive.

Finally, the comparative statics are proven by straightforward differentiation.

1.

$$\frac{\partial \Delta V}{\partial \overline{V}} = 1 - \frac{\sqrt{\overline{V} x_O W_O}}{\overline{V}} - \frac{x_F W_F}{x_F W_F + x_O W_O} > 0$$

To see this, note that $\underline{V}'_F = \sqrt{\overline{V} x_O W_O}$. Subbing in leads to

$$
\begin{aligned}
1 - \frac{\underline{V}'_F}{\overline{V}} - \frac{x_F W_F}{\underline{V}} &> 0 \\
\implies \overline{V} - \underline{V}'_F - \frac{\overline{V} x_F W_F}{\underline{V}} &> 0 \\
\implies \overline{V}\underline{V} - \underline{V}'_F\underline{V} - \overline{V} x_F W_F &> 0 \\
\implies \overline{V}\underline{V} - \overline{V} x_F W_F &> \underline{V}'_F\underline{V} \\
\implies \overline{V}(\underline{V} - x_F W_F) &> \underline{V}'_F\underline{V} \\
\implies \overline{V}(x_O W_O) &> \underline{V}'_F\underline{V} \\
\implies (\underline{V}'_F)^2 &> \underline{V}'_F\underline{V}.
\end{aligned}
$$

The final expression holds because $\underline{V}'_F > \underline{V}$.

2.

$$
\frac{\partial \Delta V}{\partial W_F} = x_F \left(1 - \frac{\overline{V} x_O W_O}{(x_F W_F + x_O W_O)^2} \right) < 0
$$

To see this, note that $\overline{V} x_O W_O = (\underline{V}'_F)^2$ and $(x_F W_F + x_O W_O)^2 = \underline{V}^2$. Subbing in leads to

$$
x_F \left(1 - \frac{(\underline{V}'_F)^2}{\underline{V}^2} \right) < 0
$$

Since $\underline{V}'_F > \underline{V}$, this expression holds.

3.

$$
\frac{\partial \Delta V}{\partial x_F} = W_F \left(1 - \frac{\overline{V} x_O W_O}{(x_F W_F + x_O W_O)^2} \right) < 0
$$

The proof is the same as 2. above.

Lemma B.1. *From the solution in Proposition 1, the following comparative statics exist:*

1. $\frac{\partial \tau^*}{\partial W_O}$ *and* $\frac{\partial \tau^*}{\partial x_O}$ *are positive if* $\overline{V} > 2\underline{V}'_F$.

2. $\frac{\partial \Delta V}{\partial W_O}$ *and* $\frac{\partial \Delta V}{\partial x_O}$ *are positive if* $(\underline{V} + \theta\overline{V}) > \frac{V\overline{V}}{\underline{V}'_F}$.

Proof 1.

$$\frac{\partial \tau^*}{\partial W_O} = x_O\left(\frac{\overline{V}}{2x_F\sqrt{\overline{V}x_O W_O}} - \frac{1}{x_F}\right)$$

This is positive if $\overline{V} > 2\underline{V}'_F$.

2.

$$\frac{\partial \tau^*}{\partial x_O} = W_O\left(\frac{\overline{V}}{2x_F\sqrt{\overline{V}x_O W_O}} - \frac{1}{x_F}\right)$$

This is positive if $\overline{V} > 2\underline{V}'_F$.

3.

$$\frac{\partial \Delta V}{\partial W_O} = x_O\left(1 - \frac{\overline{V}}{\sqrt{\overline{V}x_O W_O}} + \frac{\overline{V}x_F W_F}{(x_F W_F + x_O W_O)^2}\right)$$

This simplifies to

$$\frac{\partial \Delta V}{\partial W_O} = x_O\left(1 - \frac{\overline{V}}{\underline{V}'_F} + \frac{\theta\overline{V}}{\underline{V}}\right)$$

This is positive if $(\underline{V} - \theta\overline{V}) > \frac{V\overline{V}}{\underline{V}'_F}$

4.

$$\frac{\partial \Delta V}{\partial x_O} = W_O\left(1 - \frac{\overline{V}}{\sqrt{\overline{V}x_O W_O}} + \frac{\overline{V}x_F W_F}{(x_F W_F + x_O W_O)^2}\right)$$

This is positive if $(\underline{V} - \theta\overline{V}) > \frac{V\overline{V}}{\underline{V}'_F}$.

Lemma B.2. *Suppose that the firm reports $H_F = W_F$. Then, $V_B^O > V_N^O$, where V_N^O is the value of the project for the outsider when neither player uses a buffer and V_B^O is its value with a buffer.*

Proof Firm O will report its true cost of capital if

$$V_N^O - V_B^O = (1-\theta)(\overline{V} - \underline{V}) - \left(\overline{V} + x_F W_F - 2\sqrt{\overline{V} x_F W_F}\right) > 0.$$

Distributing and simplifying (noting that $(1-\theta)\underline{V} = X_O W_O$ from the definition if θ), the expression simplifies to:

$$-\theta\overline{V} - \underline{V} + 2\underline{V}'_O > 0$$

Solving for θ, the condition for O to not use a buffer given that O does not use a buffer is

$$\theta < \frac{2\underline{V}'_O - \underline{V}}{\overline{V}} \tag{B.8}$$

However, as shown below, B.8 can never hold. To see this start from the equivalent expression

$$\theta\overline{V} + \underline{V} - 2\underline{V}'_O < 0$$

First note that $\theta\overline{V} = \frac{x_F W_F \overline{V}}{\underline{V}}$. The numerator of this expression can be written as $\left(\sqrt{x_F W_F \overline{V}}\right)^2$. Recall that the expression $\sqrt{\overline{V} x_F W_F}$ is the new minimum acceptable project value when firm O shades and firm F does not. So, $\theta\overline{V} = \frac{(\underline{V}'_O)^2}{\underline{V}}$. We can re-write the inequality as

$$\frac{(\underline{V}'_O)^2}{\underline{V}} + \underline{V} - 2\underline{V}'_O < 0$$

$$\implies (\underline{V}'_O)^2 + \underline{V}^2 - 2\underline{V}'_O\underline{V} < 0 \tag{B.9}$$

This expression on the left side is of the form $a^2 + b^2 - 2ab$, which is strictly greater than zero for all values of a and b. Therefore, B.9 is never true and player O will always use a buffer if player F does not.

Proof of Proposition 2

Proof Integrating (3.6) over all incentive compatible gross returns ($\underline{V}''$ to $\overline{V}$) yields

$$V_F = \frac{x_F \tau_F}{x_F \tau_F + x_O \tau_O} V \Big|_{\underline{V}''}^{\overline{V}}.$$

Substituting in $\underline{V}'' = x_F \tau_F + x_O \tau_O$ yields the following:

$$V_F = \frac{x_F \tau_F}{x_F \tau_F + x_O \tau_O} [\overline{V} - (x_F \tau_F + x_O \tau_O)].$$

Simplifying gives the expression for the value of the firm in terms of τ_F:

$$V_F = \frac{x_F \tau_F \overline{V}}{x_F \tau_F + x_O \tau_O} - x_F \tau_F. \tag{B.10}$$

Similar calculation for the outside trading partner yields

$$V_O = \frac{x_O \tau_O \overline{V}}{x_F \tau_F + x_O \tau_O} - x_O \tau_O. \tag{B.11}$$

Taking first order conditions with respect to τ_F and τ_O

$$\frac{\partial V_F}{\partial \tau_F} = \frac{x_F \overline{V}(x_F \tau_F + x_O \tau_O) - x_F^2 \tau_F \overline{V}}{(x_F \tau_F + x_O \tau_O)^2} - x_F = 0$$

$$\implies \overline{V} x_O \tau_O = (x_F \tau_F + x_O \tau_O)^2, \tag{B.12}$$

$$\frac{\partial V_F}{\partial \tau_O} = \frac{x_O \overline{V}(x_F \tau_F + x_O \tau_O) - x_O^2 \tau_O \overline{V}}{(x_F \tau_F + x_O \tau_O)^2} - x_O = 0$$

$$\implies \overline{V} x_F \tau_F = (x_F \tau_F + x_O \tau_O)^2. \tag{B.13}$$

It follows that

$$\overline{V} x_F \tau_F = \overline{V} x_O \tau_O. \tag{B.14}$$

Substituting (B.14) into (B.12) for $x_O \tau_O$ gives us

$$\overline{V} x_F \tau_F = 4(x_F \tau_F)^2$$

$$\implies \tau_F^* = \frac{\overline{V}}{4x_F}.$$

By the same method

$$\overline{V} x_O \tau_O = 4(x_O \tau_O)^2$$

$$\implies \tau_O^* = \frac{\overline{V}}{4x_O}.$$

Taking second order conditions for each party yields

$$\frac{\partial^2 V_F}{\partial {\tau_F}^2} = \frac{-2x_F^2 x_O \tau_O \overline{V}(x_F \tau_F + x_O \tau_O)}{(x_F \tau_F + x_O \tau_O)^4} < 0$$

$$\frac{\partial^2 V_O}{\partial {\tau_O}^2} = \frac{-2x_O^2 x_F \tau_F \overline{V}(x_F \tau_F + x_O \tau_O)}{(x_F \tau_F + x_O \tau_O)^4} < 0,$$

so that τ_F^* and τ_O^* are global maxima.

At the optimal buffered discount rates, the value of of the project

$$V_F^* = \frac{x_F \tau_F^* \overline{V}}{x_F \tau_F^* + x_O \tau_O^*} - x_F \tau_F^*.$$

Using (B.14) to substitute in for $x_O \tau_O^*$, this simplifies to:

$$V_F^* = \frac{\overline{V}}{2} - x_F \tau_F^*.$$

Substituting (3.11) for τ_F results in

$$V_F^* = \frac{\overline{V}}{2} - x_F \frac{\overline{V}}{4x_F}$$

$$\implies = \frac{\overline{V}}{4}, \tag{B.15}$$

which is (3.12) in the text.

In what follows, we show that there is no profitable deviation for each party, if their trading partner chooses a buffer as defined above. The logic is the same for each, so we present the analysis for firm F.

If firm F *does not* use a buffer and the outsider does, the value to F is

$$V^F_{B^O} = \frac{x_F W_F}{x_F W_F + x_O \tau_O}[\overline{V} - \underline{V}'_O],$$

and the value to the firm if both use the buffer is

$$V^F_{B^{O,F}} = \frac{\overline{V}}{4}$$

Therefore, F will choose to use a buffer given O does if

$$V^F_{B^{O,F}} - V^F_{B^O} = \frac{\overline{V}}{4} - \frac{x_F W_F}{x_F W_F + x_O \tau_O}[\overline{V} - \underline{V}'_O] > 0 \tag{B.16}$$

Below we show that (B.16) is always true. First, note that $x_F W_F + x_O \tau_O$ is $\underline{V}'_O$, the minimum acceptable project when firm O uses a buffer and firm F does not. (B.16) simplifies to

$$\frac{\overline{V}}{4} - \frac{x_F W_F}{\underline{V}'_O}[\overline{V} - \underline{V}'_O] > 0$$

$$\Longrightarrow \frac{\overline{V}}{4} - \left(\frac{\overline{V} x_F W_F}{\underline{V}'_O} - x_F W_F\right) > 0 \tag{B.17}$$

The numerator of $\frac{\overline{V} x_F W_F}{\underline{V}'_O}$ can be written as $\left(\sqrt{\overline{V} x_F W_F}\right)^2$. Recall that the expression $\sqrt{\overline{V} x_F W_F}$ is the new minimum acceptable project value when firm O shades and firm F does not. So, $\frac{\overline{V} x_F W_F}{\underline{V}'_O} = \frac{(V'_O)^2}{V'_O} = V'_O$. Therefore, the expression in (B.17) simplifies to

$$\frac{\overline{V}}{4} - \underline{V}'_O + x_F W_F > 0$$

V_O' can also be written as $x_F W_F + x_O \tau_O$ which yields

$$\frac{\overline{V}}{4} - (x_F W_F + x_O \tau_O) + x_F W_F > 0$$

$$\implies \frac{\overline{V}}{4} - x_O \tau_O > 0$$

$$\implies \overline{V} > 4 x_O \tau_O \tag{B.18}$$

Next, we show that $\overline{V}$ can be expressed as a function of $x_F W_F$ and $x_O \tau_O$. The minimum acceptable project when firm O uses a buffer and F does not is defined as $x_F W_F + x_O \tau_O$. From (B.5) it follows that

$$x_F W_F + x_O \tau_O = \sqrt{\overline{V} x_F W_F}\,.$$

Squaring both sides and solving for $\overline{V}$ yields

$$(x_F W_F)^2 + 2(x_F W_F)(x_O \tau_O) + (x_O \tau_O)^2 = \overline{V} x_F W_F$$

$$\implies (x_F W_F) + 2(x_O \tau_O) + \frac{(x_O \tau_O)^2}{x_F W_F} = \overline{V}\,. \tag{B.19}$$

Substituting (B.19) into (B.18) shows that firm F will use a buffer given firm O does if

$$(x_F W_F) + 2(x_O \tau_O) + \frac{(x_O \tau_O)^2}{x_F W_F} > 4 x_O \tau_O$$

$$\implies (x_F W_F) - 2(x_O \tau_O) + \frac{(x_O \tau_O)^2}{x_F W_F} > 0$$

$$\implies (x_F W_F)^2 - 2(x_O \tau_O)(x_F W_F) + (x_O \tau_O)^2 > 0$$

This expression on the left side is of the form $a^2 + b^2 - 2ab$, which is strictly greater than zero for all values of a and b, and therefore the inequality always holds.

Finally, the projects that are undertaken when no IRR buffers are used yield an aggregate value of $\overline{V} - (x_F W_F + x_O W_O)$. The aggregate value with competitive IRR buffers is $\frac{\overline{V}}{2}$. The deadweight loss is computed as the difference.

Proof of Proposition 3

Proof The value for party i given the actions of the others is

$$V_i = \int_{\underline{R}}^{\overline{R}} \theta_i dR,$$

where $\underline{R} = \sum_{j \in N} H_j \equiv \Sigma$. This implies that

$$V_i = \frac{\overline{R} H_i}{\Sigma} - H_i. \tag{B.20}$$

Taking first-order conditions yields

$$\frac{\partial V_i}{\partial H_i} = \frac{\overline{R}\Sigma - \overline{R} H_i}{\Sigma^2} - 1. \tag{B.21}$$

Second-order conditions confirm that the objective function is strictly concave. Setting (B.21) equal to zero and re-arranging yields

$$\overline{R} \sum_{j \in J} H_j = \Sigma^2,$$

where J are all others except party i. For all pairs of parties i and j, $S \equiv \sum_{k \neq i,j}$ is the same. Therefore,

$$\overline{R}(S + H_j) = \Sigma^2 \tag{B.22}$$

$$\overline{R}(S + H_i) = \Sigma^2, \tag{B.23}$$

which implies that $H_i = H_j = H^*$ for all pairs i, j. Thus,

$$\overline{R}(n-1)H^* = n^2 (H^*)^2,$$

or

$$H^* = \frac{(n-1)\overline{R}}{n^2}. \tag{B.24}$$

Plugging (B.24) into (B.20)

$$V_i = \frac{\overline{R}}{n^2}.$$

Proof of Proposition 4

Proof With a buffer $H_F = W_F + b$, the Nash Bargaining split is

$$\theta_B = \frac{H_F}{H_F + W_O} \tag{B.25}$$

and the value of the firm can be calculated as

$$V_B = \frac{W_F\overline{R} + b\overline{R}}{W_F + b + W_O} - (W_F + b). \tag{B.26}$$

Without a buffer, the Nash Bargaining split is

$$\theta_N = \frac{W_F}{W_F + W_O} \tag{B.27}$$

and the value of the firm can be calculated as

$$V_N = \frac{W_F\overline{R}}{W_F + W_O} - W_F. \tag{B.28}$$

Therefore, $V_B > V_N$ if

$$\begin{aligned} V_B > V_N \implies \\ \frac{W_F\overline{R} + b\overline{R}}{W_F + b + W_O} - (W_F + b) > \frac{W_F\overline{R}}{W_F + W_O} - W_F \implies \\ \frac{W_F\overline{R} + b\overline{R}}{W_F + b + W_O} > \frac{W_F\overline{R}}{W_F + W_O} + b \implies \\ \frac{\overline{R}W_O - (W_F + W_O)^2}{(W_F + W_O)} > b \end{aligned} \tag{B.29}$$

Proof of Proposition 5

Proof The hurdle rate, H_F, that implements the CFO's limitation on feasible projects is set such that the following holds with equality:

$$k[\overline{R} - W_F - W_O] = \overline{R} - H_F - W_O$$

Solving for H_F yields

$$H_F = (1 - k)\overline{R} + kW_F - (1 - k)W_O \Longrightarrow$$

$$H_F = kW_F + (1 - k)(\overline{R} - W_O). \tag{B.30}$$

The value of the project under the biased hurdle rate is:

$$V_B = \int_{\underline{R}_H}^{\overline{R}} \theta_H dR = \theta_H[\overline{R} - (H_F + W_O)] = \frac{H_F}{H_F + W_O} - H_F$$

The value under the biased hurdle rate exceeds the value under the true discount rate if:

$$V_B > V_N \Longrightarrow$$

$$\theta_H[\overline{R} - (H_F + W_O)] > \theta_N[\overline{R} - (W_F + W_O)] \Longrightarrow$$

$$\theta_H[\overline{R} - \underline{R}_H] > \theta_N[\overline{R} - \underline{R}] \Longrightarrow$$

$$(\theta_H - \theta_N)\overline{R} > (H_F - W_F) \Longrightarrow$$

$$(\theta_H - \theta_N) > (1 - k)(1 - \frac{W_F + W_O}{\overline{R}}). \tag{B.31}$$

The last inequality follows from the fact that

$$H_F - W_F = (1 - k)[\overline{R} - (W_F + W_O)].$$

The optimal hurdle rate that maximized the firm's share of the surplus, τ^*, is greater than the

biased hurdle rate, H_F, if:

$$\begin{aligned}
\tau^* > H_F \implies \\
\sqrt{\overline{R}W_O} - W_O > (1-k)\overline{R} + k(W_F + W_O) - W_O \implies \\
\underline{R}' > (1-k)\overline{R} + k\underline{R} \implies \\
k[\overline{R} - \underline{R}] > \overline{R} - \underline{R}' \implies \\
k > \frac{\overline{R} - \underline{R}'}{\overline{R} - \underline{R}}
\end{aligned} \tag{B.32}$$

B.3 Data Appendix

The CFO survey has asked respondents about their hurdle rate and cost of capital jointly six different times: 2011q1, 2012q2, 2017q2, 2017q3, 2019q1 and 2022q2. The last survey was conducted with the Federal Reserve Banks of Richmond and Atlanta. On each survey, the wording of the hurdle rate question specifically asked CFOs for their investment hurdle rates (the expected rate of return an investment project must exceed in order to be adopted). We similarly asked CFOs to supply their weighted average cost of capital (WACC), as opposed to their cost of equity or debt. As an example, dispayed below is asked CFOs for their hurdle rate and WACC in the 2012q2 survey. The questions for the other surveys can be found at https://cfosurvey.fuqua.duke.edu/ for surveys conducted prior to 2022q2 and https://www.richmondfed.org/cfosurvey/ for the 2022q2 survey.

10a. What do you estimate is your firm's overall weighted average cost of capital (WACC)? (e.g., 11.2%)	10b. What is your firm's 'hurdle rate' (the rate of return that an investment must beat in order to be adopted)?
%	%

A potential concern with survey data is that respondents do not understand the survey questions, or do not respond accurately (Graham, 2022). While we cannot fully address these concerns, we can analyze the accuracy of the CFO survey forecasts for firms that also have archival data in Compustat (about 15% of the sample). Following analysis in Gormsen and Huber (2020) that relates firms' survey costs of capital to estimated costs of capital, Figure B.5 Panel A displays a binned scatter plot of the WACC as reported on the survey and the WACC we estimate using data from CRSP and Compustat. We estimate WACC in the simplest way possible. The cost of equity is derived using the CAPM; the cost of debt is total interest expense to total debt (i.e., the average cost of debt). The slope coefficient from this regression is 0.61 and highly significant, and the R^2 is 0.44. While exploring the relation between perceived and estimated costs of capital is not the point of our paper, we are reassured by the tightness of the relation. Secondly, Figure B.5 Panel B compares the survey-reported revenue categories of these same firms to the equivalent category calculated using Compustat; the proportions line up nearly one-to-one. Finally, recall the analysis of Figure 3.4, which indicates that ex post realized returns align with survey-provided hurdle rates, again providing assurance that our survey data are reliable.

A second concern with survey data is representativeness. Table IAI in the Internet Appendix of Graham (2022), available at https://onlinelibrary.wiley.com/doi/10.1111/jofi.13161, shows that, even though the CFO survey contains a large proportion of private firms, the distributions of

employment counts conditional on firm size are quite similar in the CFO survey and Compustat. Thus, we do not detect anomalies that would suggest problems using our data to draw inference relative to well-known samples such as Compustat.

Appendix C

Appendix to Chapter 4

C.1 Appendix Figures and Tables

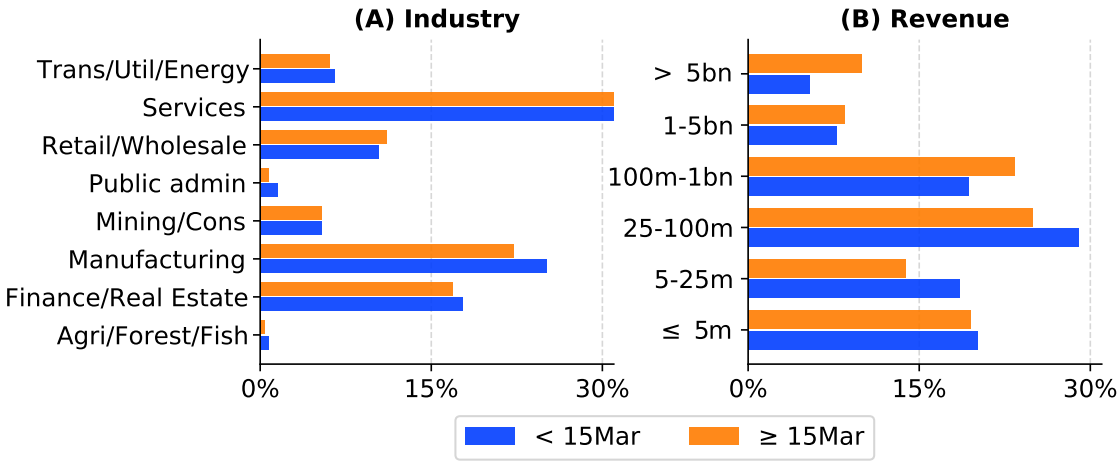

FIGURE C.1: Sample Composition by March 2020 Survey Completion Date

This figure displays the composition of firms in the March 2020 CFO survey split by pre/post March 15, by industry (Panel A) and firm size (Panel B).

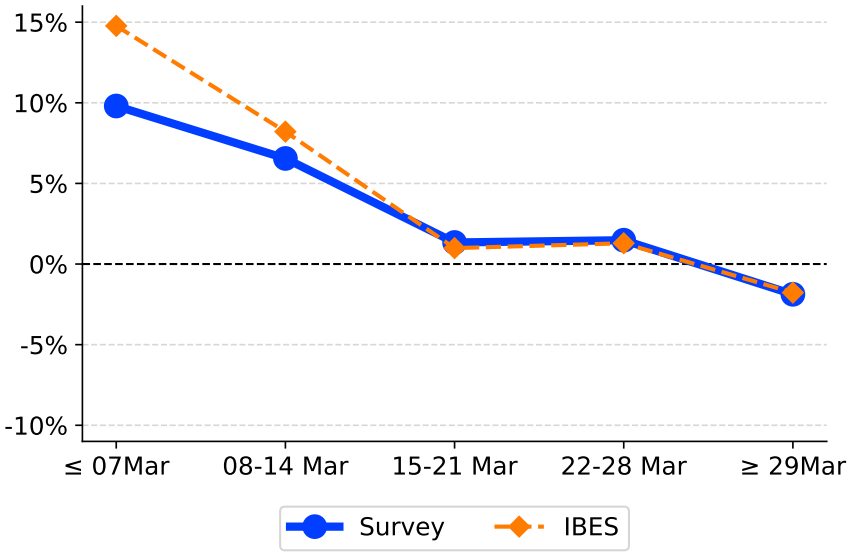

FIGURE C.2: CFO and IBES Forecasts of 2020 Revenue Growth

The solid line shows the average CFO forecast of revenue growth in 2020 by survey week. The data come from the March 2020 CFO survey. The dashed line shows the contemporaneous average analyst forecast of revenue growth in 2020 from IBES among analysts who issue new forecasts each week (we do not include past forecasts that were not updated in a given week).

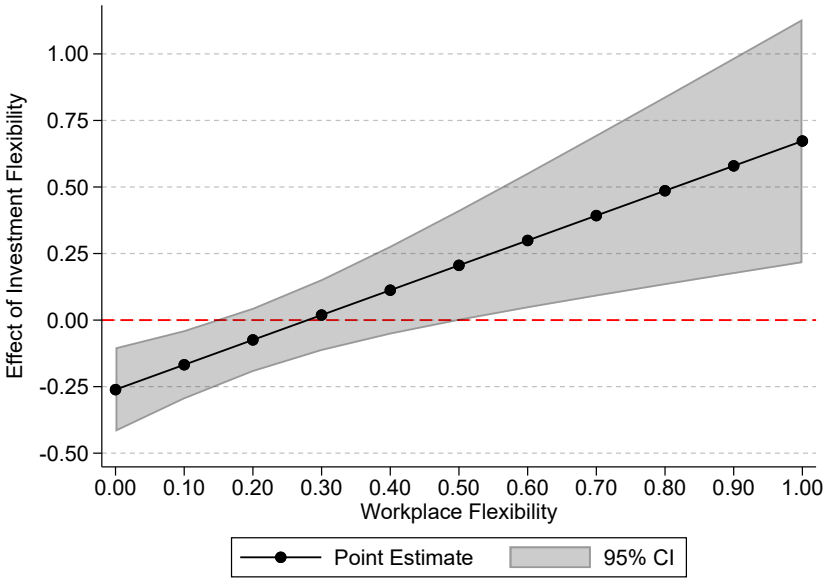

FIGURE C.3: Effect of Investment Flexibility on Capital Spending Forecasts for Different Levels of Workplace Flexibility

This figure displays the effect of investment flexibility on capital spending forecasts, over the range of workplace flexibility. Estimated using column (5) of Table 4.3, the estimating equation is

$$\begin{aligned}\text{Capital Spending Forecast}_{it} = \alpha + \beta_1 \text{Investment Flex}_{it} + \beta_2 \text{Workplace Flex}_{it} + \\ \beta_3(\text{Investment Flex} \times \text{Workplace Flex}) + \lambda \cdot X_{it} + \varepsilon_{it}\end{aligned}$$

Each point on the black line displays the average marginal effect of investment flexibility on capital spending forecasts, for a given value of workplace flexibility, $\mathbb{E}[\text{Marginal Efect} \mid \text{Workplace Flex} = w] = \beta_1 + \beta_3 w$. The shaded area displays 95% confidence intervals. The data come from the March 2020 survey and use CFOs' projected capital spending growth in 2020.

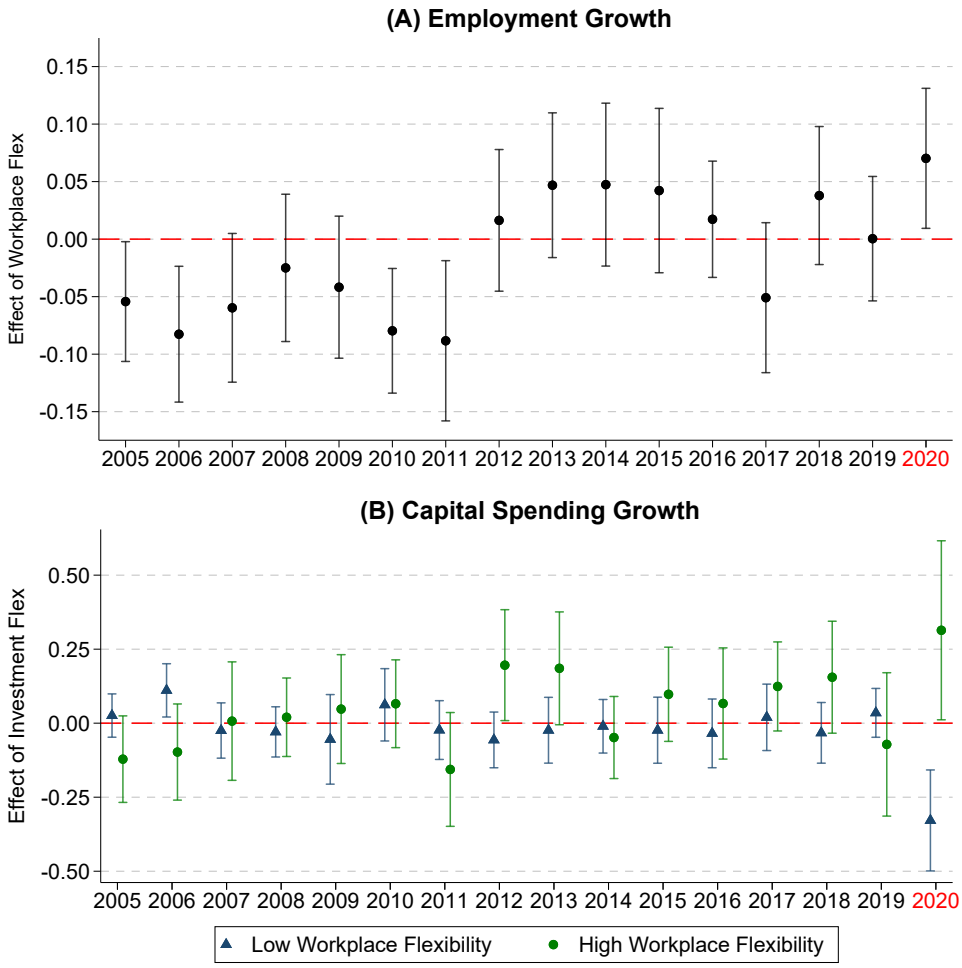

FIGURE C.4: Employment and Capital Spending Growth in Compustat in Prior Years

This figure displays the effect of workplace flexibility and investment flexibility on realized employment and capital spending growth throughout time. In Panel A, we plot the coefficient on Workplace Flexibility from annual cross-sectional specifications where annual realized employment growth is regressed on Workplace Flexibility, Investment Flexibility, along with controls (lagged leverage, lagged cash/assets, lagged log # employees) and fixed effects (NAICS-2 and state). The exact specification mirrors that in Table 4.7, column (3). In Panel B, we set the sample as all eligible Compustat observations for each year. We then regress log quarterly capital spending growth on the interaction of workplace and investment flexibility, along with controls and fixed effects (which are the same as described in Panel A above). The exact specification mirrors that in Table 4.7, column (6). The blue triangles (green dots) display the effect of investment flexibility on realized capital spending growth for a firm with low (high) workplace flexibility. Low (high) workplace flexibility is defined as the 10th (90th) percentile value of workplace flexibility at the four-digit NAICS level (0.036 and 0.653, respectively). In both panels, the vertical bars display 95% confidence intervals.

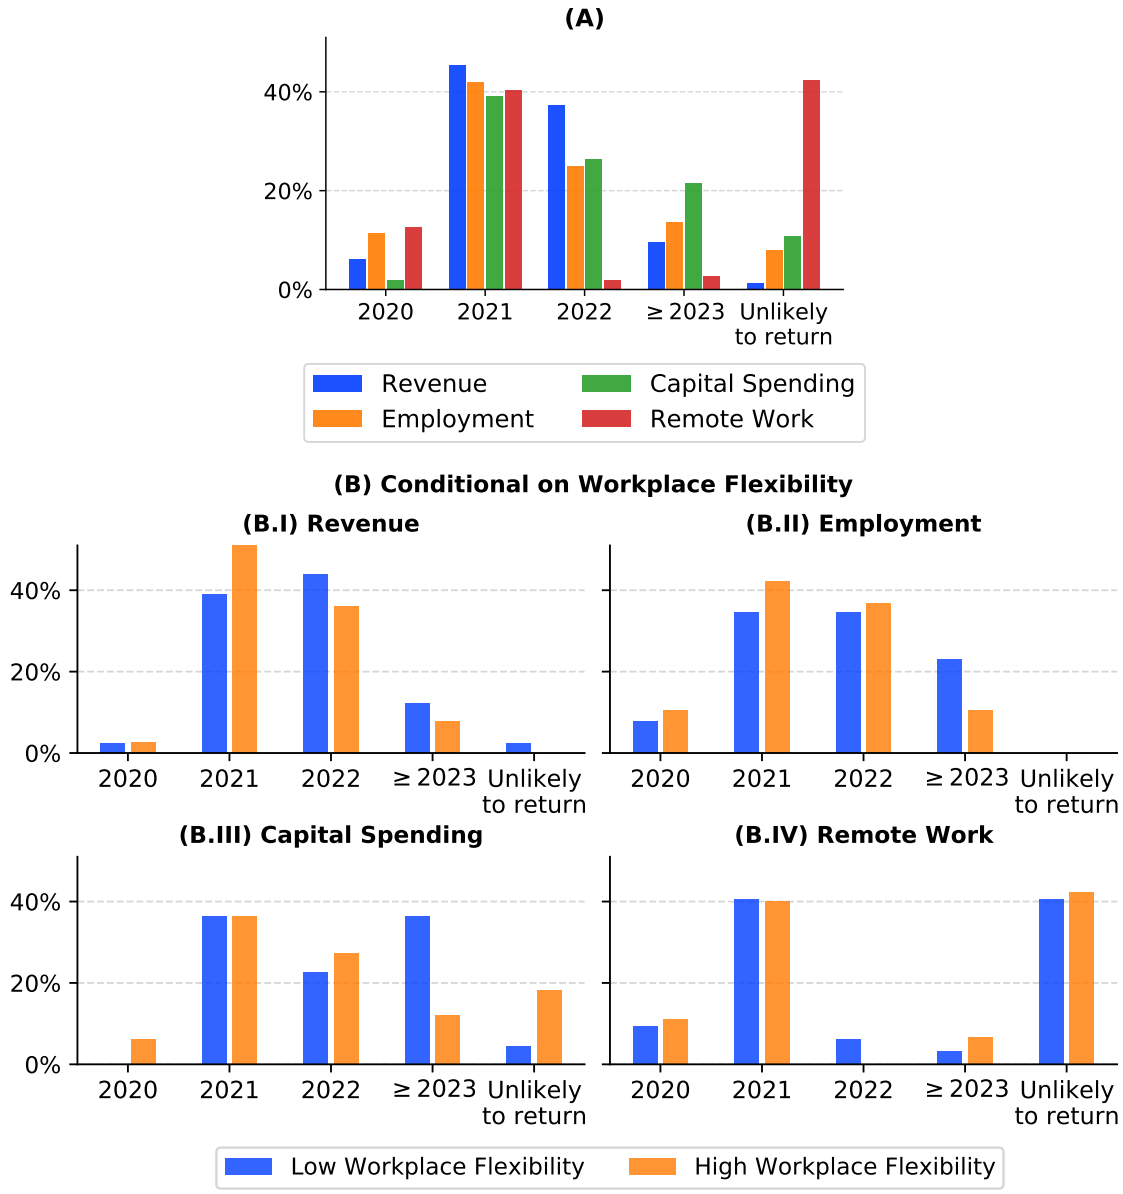

FIGURE C.5: CFO Outlook of Firm Outcomes Returning to Pre-COVID Levels

Each panel displays the time period at which CFOs expect the relevant variable to return to pre-COVID levels. Revenue, employment and remote work refer to the level of the variable. Capital spending refers to "willingness to spend on structures and equipment." For the revenue, employment and capital spending panels, the sample is limited to firms that stated they saw a decrease in the relevant variable since the onset of COVID-19. For the remote work panel, the sample is limited to firms that stated they saw an increase in remote work since the onset of COVID-19. CFOs reporting no change to the relevant variable are omitted from the calculations. Data are from the September 2020 CFO survey. Panel A is for all firms that stated they saw a decrease in revenue, employment, or capital spending, or an increase in remote work. Panel B displays by workplace flexibility. Low (high) workplace flexibility is below (above) the 25th (75th) percentile of workplace flexibility within-sample.

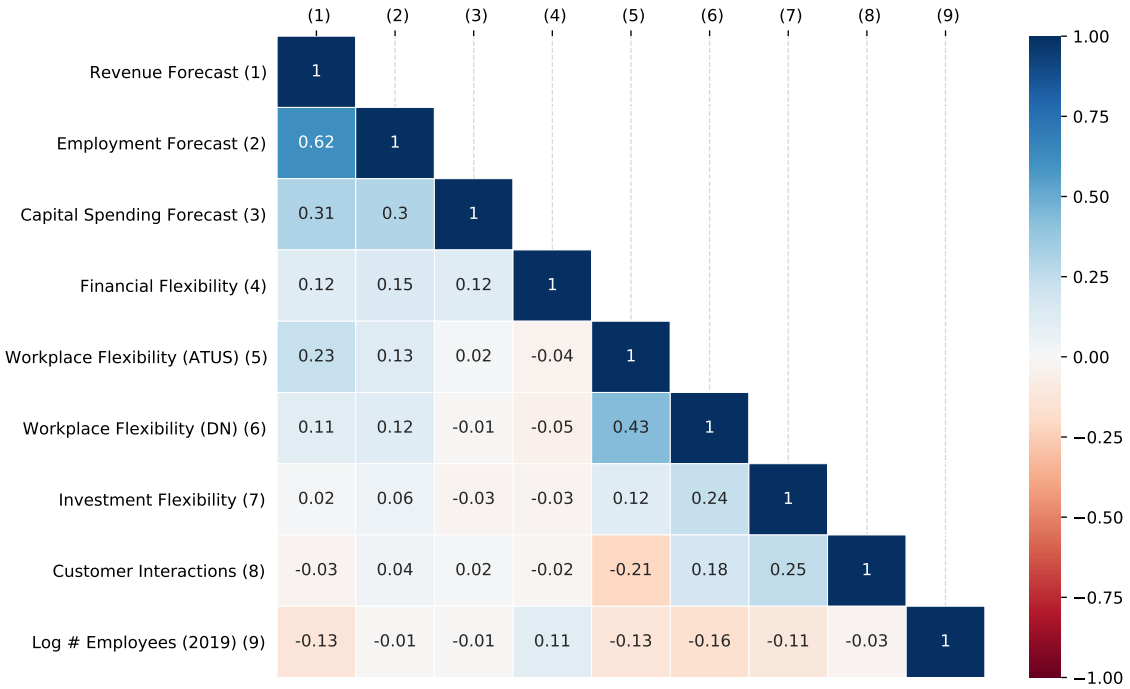

FIGURE C.6: Cross-Correlations from March 2020 Survey

This figure shows the correlations among the main variables. Dark blue indicates strong positive correlations, and dark red indicates strong negative correlations. Data are from the March 2020 CFO survey. Detailed variable definitions are in Appendix C.3.

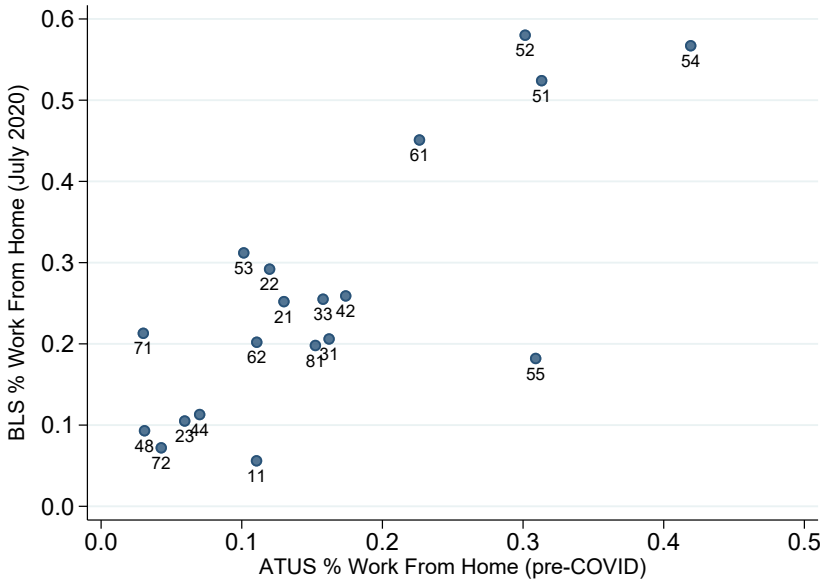

FIGURE C.7: Remote Work in ATUS (Pre-COVID) versus BLS Data (July 2020)

This figure shows the correlation between work from home measured in ATUS and in BLS data. Each point is a two-digit NAICS industry.

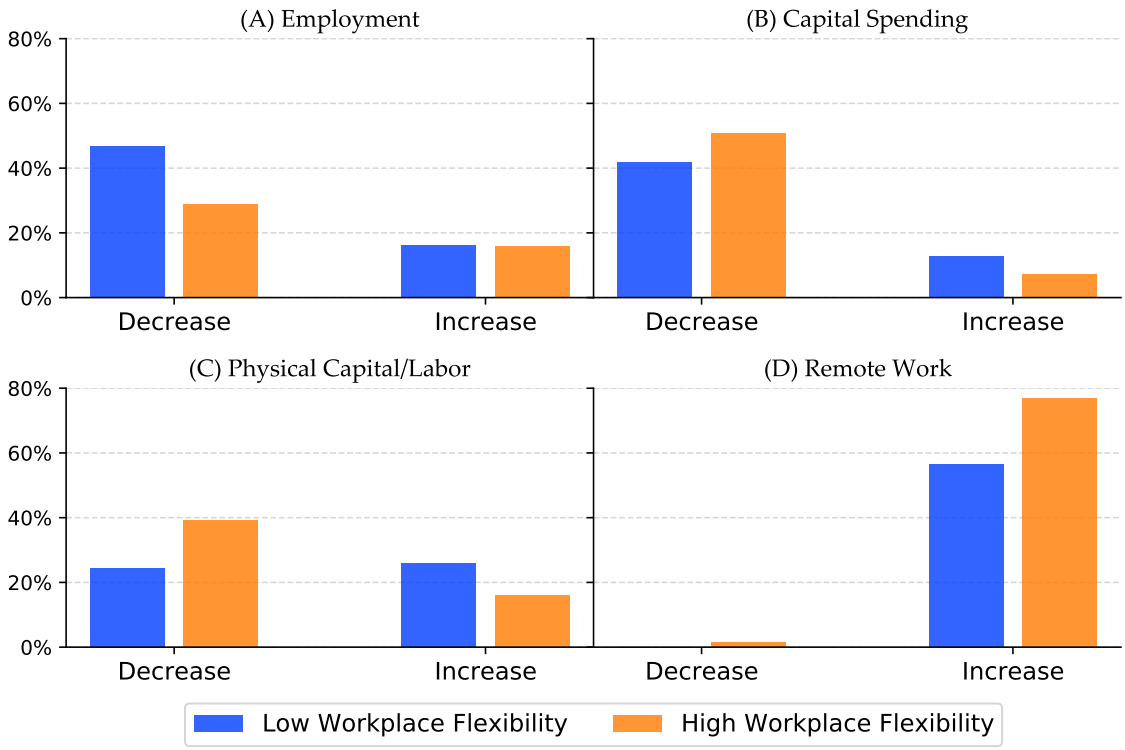

FIGURE C.8: Effect of COVID-19 on Company Outcomes

Each panel displays the percentage of CFOs who stated that the current level for their company was lower (higher) than pre-COVID level. The data come from the September 2020 CFO survey. Employment, capital spending and remote work refer directly to whether the level of the variable decreased or increased in September 2020, relative to pre-COVID. For example, the orange bar above "Decrease" in Panel A indicates that about 30% of high workplace flexibility firms had reduced employment as of September 2020; the neighboring blue bar indicates that about 45% of low workplace flexibility firms had reduced employment as of September 2020. "Physical Capital/Labor" is coded as "Decrease" ("Increase") if the new level of capital spending is lower (higher) than the new level of employment (refer to Table C.10 for a detailed definition). Capital spending refers to "willingness to spend" on structures and equipment. Firms that stated there has been no change are omitted from the figure, thus within-group bars do not sum to one. Low (high) workplace flexibility is below (above) the 25th (75th) percentile of workplace flexibility within-sample.

Table C.1: Determinants of Financial Flexibility

	(1)	(2)	(3)
	Linear Probability Model		Logit
Cash/Assets	0.285***	0.512***	0.348***
	(0.066)	(0.089)	(0.095)
Limited Access to External Capital	-0.169***	-0.163***	-0.167***
	(0.034)	(0.025)	(0.040)
Observations	454	448	454
R-squared	0.060	0.212	0.063
Week FE		Yes	
State FE		Yes	
NAICS-2 FE		Yes	

This table examines the determinants of firms' self-assessed financial flexibility in the March 2020 survey. In all specifications, the dependent variable is an indicator variable taking a value of one if the firm stated they had more financial flexibility than "None" or "A little." Columns (1) and (2) present results from linear probability models (OLS), column (3) presents results from a logit specification. Cash/Assets is the firm's stated cash to total assets ratio from year-end 2019. Limited Access to External Capital is an indicator taking a value of one if the firm stated that their ability to access external capital limited their ability to pursue attractive investment projects. Detailed variable definitions are available in Appendix C.3. The R-squared in column (3) is the pseudo R-squared from the logit regression. Standard errors are clustered at the two-digit NAICS level and displayed in parentheses below the coefficient. ***, **, * denote significance at 1%, 5%, 10%.

Table C.2: Investment Flexibility by 2-Digit NAICS Industry

NAICS-2	Description	Investment Flexibility
81	Other Services (except Public Administration)	0.667
23	Construction	0.600
71	Arts, Entertainment, and Recreation	0.500
51	Information	0.420
52	Finance and Insurance	0.411
61	Educational Services	0.333
44-45	Retail Trade	0.300
54	Professional, Scientific, and Technical Services	0.269
56	Administrative and Support and Waste Managemen...	0.250
42	Wholesale Trade	0.175
31-33	Manufacturing	0.151
21	Mining, Quarrying, and Oil and Gas Extraction	0.125
48-49	Transportation and Warehousing	0.125
53	Real Estate and Rental and Leasing	0.100
62	Health Care and Social Assistance	0.050
11	Agriculture, Forestry, Fishing and Hunting	0
22	Utilities	0
72	Accommodation and Food Services	0

This table displays our measure of investment flexibility at the 2-digit NAICS level. This broader industry measure is illustrative - the analysis of the paper uses the finer 4-digit NAICS measure, as defined in Appendix C.3. The 4-digit data and details on construction can be found at https://github.com/jwb4335/corporate_flexibility.

Table C.3: Firm-Specific Measure of Investment Flexibility from March 2019 Subsample

	(1)	(2)	(3)	(4)	(5)	(6)
	Employment			Capital Spending		
Financial Flexibility	0.073**	0.070*	0.092*	0.039	0.349*	0.258
	(0.028)	(0.039)	(0.043)	(0.055)	(0.175)	(0.199)
Workplace Flexibility	0.111**	0.078**	0.085**	-0.250**	-0.293	-0.318
	(0.042)	(0.035)	(0.034)	(0.105)	(0.186)	(0.189)
Investment Flexibility	-0.032	-0.007	0.078	-0.166	-0.307***	-0.736**
	(0.030)	(0.026)	(0.045)	(0.117)	(0.084)	(0.273)
Workplace Flex × Investment Flex	0.015	-0.017	-0.033	0.613**	0.950**	0.999**
	(0.043)	(0.044)	(0.042)	(0.270)	(0.326)	(0.344)
Financial Flex × Investment Flex			-0.088			0.443
			(0.059)			(0.278)
Observations	145	132	132	140	128	128
R-squared	0.083	0.488	0.494	0.026	0.484	0.499
Controls		Yes	Yes		Yes	Yes
Week FE		Yes	Yes		Yes	Yes
State FE		Yes	Yes		Yes	Yes
NAICS-2 FE		Yes	Yes		Yes	Yes

This table examines the interactive effects of workplace and investment flexibility on firms' employment and capital spending plans for firms with a direct measure of investment flexibility, in support of Table 4.3. The sample is firms that appeared in both the March 2019 (and supplied a direct measure of their firm's Investment Flexibility) and March 2020 CFO Surveys. The dependent variable is the CFOs' projected growth rate for employment (columns (1) to (3)) or capital spending (columns (4) to (6)) for 2020, as given in the March 2020 survey. Investment Flexibility is a dummy variable taking a value of one if the CFO stated they "Flexible" or "Very Flexible" speed with respect to their (largest) investment projects, i.e. the base variable we use to construct our industry-level measure of Investment Flexibility used in our main tests. Financial and Workplace Flexibility are as defined in Table 4.3. Controls are Customer Interactions at Log # Employees (at the end of 2019). Detailed variable definitions are available in the Appendix C.3. Standard errors are clustered at the two-digit NAICS level and displayed in parentheses below the coefficient. ***, **, * denote significance at 1%, 5%, 10%.

Table C.4: Realized Outcomes in Compustat Data: Impact of Fixed Costs

	(1)	(2)	(3)	(4)	(5)	(6)
	Employment Growth					
Fixed Cost Share	-0.204***	-0.145**	-0.145**	-0.722***	-0.613***	-0.608***
	(0.066)	(0.066)	(0.061)	(0.147)	(0.187)	(0.183)
Financial Flexibility Proxy	0.067	0.048	0.034			
	(0.063)	(0.073)	(0.071)			
Fixed Cost Share × Fin Flex Proxy	0.394***	0.390***	0.382***			
	(0.093)	(0.111)	(0.111)			
Financial Flexibility Proxy (Fitted)				-0.052	-0.058	-0.086
				(0.089)	(0.112)	(0.107)
Fixed Cost Share × Fin Flex Proxy (Fitted)				0.756***	0.694***	0.681***
				(0.141)	(0.198)	(0.200)
Observations	4,574	4,572	4,572	4,574	4,572	4,572
R-squared	0.064	0.097	0.101	0.055	0.089	0.093
State FE		Yes	Yes		Yes	Yes
NAICS-2 FE		Yes	Yes		Yes	Yes
Controls			Yes			Yes

This table provides external validation from Compustat for our results concerning the interaction of financial flexibility and a firm's fixed cost share found in Table 4.4. We use Compustat data from fiscal year 2020. We require that a firm have positive assets, non-negative debt, non-missing data for lagged leverage and cash/assets, lagged leverage between zero and one, non-missing employment data from fiscal years 2019 and 2020, and a non-missing four-digit NAICS code. We consider two proxies for a firm's financial flexibility. In columns (1) to (3), Financial Flexibility Proxy is the simple average of the firm's lagged cash/assets (CHE/AT) and one minus the firm's lagged leverage (1-(DLC+DLTT)/AT). In columns (4) to (6), we use evidence on the determinants of financial flexibility found in Table C.1 to develop a fitted measure of financial flexibility. Specifically, Financial Flexibility Proxy (Fitted) uses the coefficients on Cash/Assets and the indicator variable Limited Access to External Capital in Table C.1, column (1), along with each Compustat firm's lagged cash/assets and an indicator taking a value of one if the firm's lagged leverage is below median (i.e., above median leverage proxies for limited access to external capital). Controls are Customer Interactions, as defined in Appendix C.3, and Lagged Log # Employees, the natural logarithm of the firm's number of employees from the previous fiscal year (2019). Standard errors are clustered at the four-digit NAICS level and displayed in parentheses below the coefficient. ***, **, * denote significance at 1%, 5%, 10%.

Table C.5: Descriptive Statistics from Subsequent Surveys

	N	Mean	Std dev	25%	Median	75%
CFO Forecast Variables						
Revenue Forecast (for 2020)	626	0.024	0.367	-0.100	0	0.100
Revenue Forecast (for 2021)	621	0.101	0.208	0.010	0.050	0.150
Employment Forecast (for 2020)	640	-0.006	0.176	-0.075	0	0.042
Employment Forecast (for 2021)	641	0.057	0.217	-0.042	0	0.125
Flexibility Variables						
Workplace Flexibility (ATUS)	641	0.214	0.157	0.065	0.224	0.334
Workplace Flexibility (DN)	641	0.470	0.261	0.225	0.418	0.762
Investment Flexibility	641	0.261	0.304	0	0.200	0.364
Control Variables						
Customer Interactions	641	0.466	0.101	0.403	0.487	0.515
Log # Employees (2019)	641	4.774	2.240	3.296	4.605	5.943

This table presents summary statistics of the main variables from the CFO surveys in June, September and December 2020. The number of observations, means, standard deviations and quartiles are displayed. CFO forecasts of revenue and employment represent growth from the end of 2019 to the end of 2020 and 2021. Detailed variable definitions are given in Appendix C.3.

Table C.6: Determinants of COVID Risk Exposure

	(1)	(2)	(3)	(4)
	Linear Probability Model			Logit
Financial Flexibility	-0.036	-0.042	-0.039	-0.036
	(0.062)	(0.054)	(0.051)	(0.052)
Workplace Flexibility	-0.227**	-0.182**	-0.181**	-0.245***
	(0.088)	(0.081)	(0.075)	(0.083)
Investment Flexibility	-0.118	-0.145*	-0.139*	-0.114*
	(0.070)	(0.073)	(0.073)	(0.061)
Customer Interactions	0.498***	0.502***	0.549***	0.509***
	(0.082)	(0.103)	(0.090)	(0.073)
Log # Employees (2019)	0.010	0.013	0.013	0.010
	(0.011)	(0.011)	(0.011)	(0.010)
Post March 15	0.358***	0.377***		0.358***
	(0.056)	(0.061)		(0.050)
Observations	451	445	445	451
R-squared	0.166	0.268	0.291	0.127
State FE		Yes	Yes	
Week FE			Yes	

This table examines the determinants of firms' self-assessed exposure to COVID risk. In all specifications, the dependent variable is an indicator variable taking a value of one if firms in the March 2020 survey stated they faced medium or high coronavirus risk. Columns (1) to (3) present results from linear probability models (OLS), and column (4) presents results from a logit specification. The dependent variable in all specifications is an indicator variable equal to one if the CFO stated their firm faced "medium" or "large" Coronavirus risk and zero otherwise. Financial Flexibility is an indicator taking a value of one if the firm stated they had more financial flexibility than "None" or "A little." Workplace Flexibility comes from ATUS and is a four-digit NAICS level measure for the percentage of workers that can work from home. Investment Flexibility is a four-digit NAICS level measure for a firm's investment flexibility (with respect to speed of completion). Customer Interactions is a four-digit NAICS level variable that proxies for the intensity of interactions with consumers. Log # Employees (2019) is the natural logarithm of the firm's number of employees at the end of 2019. Detailed variable definitions are available in Appendix C.3. The R-squared in column (4) is the pseudo R-squared from the logit regression. Standard errors are clustered at the two-digit NAICS level and displayed in parentheses below the coefficient. ***, **, * denote significance at 1%, 5%, 10%.

Table C.7: Conditional Impact of Investment Flexibility: Full Set of Interactions

Panel A: Impact on Capital Spending

	(1)	(2)	(3)	(4)	(5)	(6)
	Capital Spending					
Workplace Flexibility	-0.123	-0.140*	-0.124	-0.143**	-0.222	-0.329**
	(0.074)	(0.068)	(0.074)	(0.068)	(0.162)	(0.154)
Investment Flexibility	-0.197**	-0.209**	-0.373**	-0.309***	-0.363**	-0.291**
	(0.075)	(0.084)	(0.134)	(0.093)	(0.144)	(0.107)
Financial Flexibility		0.091***	0.025	0.059***	-0.001	0.005
		(0.025)	(0.039)	(0.020)	(0.044)	(0.032)
Workplace Flex × Investment Flex	0.736***	0.744**	0.769***	0.765**	0.779***	0.784***
	(0.183)	(0.278)	(0.181)	(0.278)	(0.179)	(0.272)
Financial Flex × Investment Flex			0.210	0.116	0.196	0.089
			(0.147)	(0.094)	(0.158)	(0.110)
Workplace Flex × Financial Flex					0.120	0.229
					(0.117)	(0.133)
Observations	397	391	397	391	397	391
R-squared	0.020	0.197	0.034	0.198	0.035	0.201
Controls		Yes		Yes		Yes
Week FE		Yes		Yes		Yes
State FE		Yes		Yes		Yes
NAICS-2 FE		Yes		Yes		Yes

Panel B: Impact on Employment

	(1)	(2)	(3)	(4)	(5)	(6)
	Employment					
Financial Flexibility	0.071***	0.077***	0.037*	0.048	0.027	0.039
	(0.018)	(0.019)	(0.020)	(0.034)	(0.025)	(0.045)
Workplace Flexibility	0.039**	0.046**	0.040**	0.043**	0.001	0.014
	(0.017)	(0.017)	(0.017)	(0.017)	(0.035)	(0.058)
Investment Flexibility	-0.037*	-0.025	-0.139***	-0.114*	-0.135***	-0.111*
	(0.021)	(0.027)	(0.032)	(0.057)	(0.031)	(0.054)
Workplace Flex × Investment Flex	0.302***	0.268	0.310***	0.287	0.316***	0.289
	(0.090)	(0.182)	(0.086)	(0.180)	(0.085)	(0.182)
Financial Flex × Investment Flex			0.122***	0.104*	0.116***	0.100**
			(0.033)	(0.050)	(0.032)	(0.047)
Workplace Flex × Financial Flex					0.047	0.036
					(0.034)	(0.059)
Observations	405	400	405	400	405	400
R-squared	0.058	0.230	0.065	0.234	0.065	0.234
Controls		Yes		Yes		Yes
Week FE		Yes		Yes		Yes
State FE		Yes		Yes		Yes
NAICS-2 FE		Yes		Yes		Yes

This table is an extension of Table 4.3 in the main text, where we include pairwise interactions among all three flexibility measures. Data are from the March 2020 CFO survey. The dependent variable is the projected capital spending growth in 2020 in Panel A, and the projected employment growth in 2020 in Panel B. Controls are Customer Interactions and Log # Employees (at the end of 2019). Detailed variable definitions are available in Appendix C.3. Standard errors are clustered at the two-digit NAICS level and displayed in parentheses below the coefficient. ***, **, * denote significance at 1%, 5%, 10%.

Table C.8: Investment Flexibility and Customer Interactions

Panel A: March 2020 CFO Survey

	(1)	(2)	(3)	(4)
	Capital Spending Growth			
Customer Interactions	0.491	0.066	0.002	0.025
	(0.373)	(0.345)	(0.343)	(0.348)
Investment Flexibility	0.456	0.268	0.260	-0.047
	(0.366)	(0.179)	(0.190)	(0.200)
Workplace Flexibility	0.118	0.124	0.113	-0.071
	(0.113)	(0.101)	(0.102)	(0.126)
Customer Interactions × Investment Flex	-1.041	-0.782**	-0.774**	-0.605*
	(0.678)	(0.350)	(0.374)	(0.342)
Workplace Flex × Investment Flex				0.829***
				(0.267)
Observations	4,215	4,215	4,212	4,212
R-squared	0.021	0.044	0.058	0.059
Controls	Yes	Yes	Yes	Yes
NAICS-2 FE		Yes	Yes	Yes
State FE			Yes	Yes

Panel B: Compustat Annual 2020

	(1)	(2)	(3)	(4)
	Capital Spending Growth			
Customer Interactions	0.491	0.066	0.002	0.025
	(0.373)	(0.345)	(0.343)	(0.348)
Investment Flexibility	0.456	0.268	0.260	-0.047
	(0.366)	(0.179)	(0.190)	(0.200)
Workplace Flexibility	0.118	0.124	0.113	-0.071
	(0.113)	(0.101)	(0.102)	(0.126)
Customer Interactions × Investment Flex	-1.041	-0.782**	-0.774**	-0.605*
	(0.678)	(0.350)	(0.374)	(0.342)
Workplace Flex × Investment Flex				0.829***
				(0.267)
Observations	4,215	4,215	4,212	4,212
R-squared	0.021	0.044	0.058	0.059
Controls	Yes	Yes	Yes	Yes
NAICS-2 FE		Yes	Yes	Yes
State FE			Yes	Yes

This table examines the interactive role of investment flexibility and customer interactions in determining capital spending plans and outcomes. Panel A considers capital spending plans for firms in the March 2020 CFO Survey and Panel B considers annual capital spending growth outcomes for the fiscal year 2020 for Compustat firms. In Panel A, the dependent variable is the firm's expected growth in capital spending from the end of 2019 to the end of 2020. Controls are Financial Flexibility and Log # Employees (at the end of 2019). Workplace Flexibility, Customer Interactions and Investment Flexibility are as defined in previous tables. In Panel B, the dependent variable is the log change in capital spending from 2019 to 2020. Controls are Lagged Cash/Assets, Lagged Leverage and Log # Employees (at the end of 2019). For Panel B, we require that a firm have positive assets, non-negative debt, non-missing data for lagged leverage and cash/assets, non-missing employment data from fiscal years 2019 and 2020, and a non-missing four-digit NAICS code. Standard errors are clustered at the four-digit NAICS level and displayed in parentheses below the coefficient. ***, **, * denote significance at 1%, 5%, 10%.

Table C.9: Impact of Workplace Flexibility on Employment Growth Realizations 2005-2019

	(1)	(2)	(3)	(4)	(5)	(6)
	Compustat				BLS	
Workplace Flexibility	0.051***	0.001	-0.008	-0.005	0.014	0.019
	(0.010)	(0.014)	(0.017)	(0.013)	(0.010)	(0.012)
Lagged Leverage		-0.015***	-0.018***	-0.018***		
		(0.002)	(0.002)	(0.002)		
Lagged Cash/Assets		0.098***	0.076***	0.083***		
		(0.012)	(0.014)	(0.012)		
Customer Interactions			0.054***	-0.005		
			(0.015)	(0.030)		
Lagged Log # Employees			-0.006***	-0.006***		
			(0.001)	(0.001)		
Observations	69,249	69,249	69,249	69,249	3,360	3,360
R-squared	0.001	0.009	0.022	0.024	0.002	0.314
Year FE			Yes	Yes		Yes
State FE			Yes	Yes		
NAICS-2 FE				Yes		Yes

This table examines the effect of workplace flexibility on employment growth realizations. Columns (1) to (4) display specifications relating workplace flexibility and employment growth for Compustat firms for the fiscal years 2005-2019. We require that the firm have positive assets, non-negative debt, non-missing data for lagged leverage and cash/assets, non-missing employment data from the current and previous year, and a non-missing four-digit NAICS code. Columns (5) and (6) display specifications relating workplace flexibility and employment growth at the industry level (four-digit NAICS) using data on employment counts from the Bureau of Labor Statistics (BLS) National Current Employment Statistics Survey. The dependent variable in all specifications is the log change in employment from the previous year. BLS employment growth is measured from December to December. Controls in columns (1) to (4) are Customer Interactions, as defined in Appendix C.3, and Lagged Log # Employees, the natural logarithm of the firm's number of employees from the previous fiscal year. Standard errors are clustered at the four-digit NAICS level and displayed in parentheses below the coefficient. ***, **, * denote significance at 1%, 5%, 10%.

Table C.10: Realized Company Outcomes Relative to Pre-COVID

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Employment		Capital Spending		Physical Capital/Labor		Remote Work	
Workplace Flexibility	1.317*	1.221	0.778**	0.760**	0.698***	0.707***	1.511***	2.291***
	(0.149)	(0.142)	(0.117)	(0.129)	(0.109)	(0.105)	(0.116)	(0.260)
Investment Flexibility	0.968	0.960	0.984	0.948	0.980	0.975	0.782**	0.797
	(0.181)	(0.186)	(0.090)	(0.103)	(0.140)	(0.149)	(0.105)	(0.139)
Customer Interactions		0.841**		1.095		1.275**		1.107
		(0.077)		(0.108)		(0.102)		(0.221)
Log # Employees (2019)		0.604***		0.914		1.491***		4.237***
		(0.110)		(0.165)		(0.150)		(0.267)
Observations	244	244	244	244	244	244	244	244
Pseudo R-squared	0.010	0.046	0.012	0.014	0.022	0.048	0.033	0.208

This table examines how firm outcomes have changed since the onset of COVID, using the same data as Figure C.8. Data are from the September 2020 CFO survey. This survey asked CFOs:
For your company, how would you assess the current level of {Employment, Capital Expenditure (Willingness to Spend on Structures and Equipment), Remote Work} compared to their levels before the outbreak of COVID-19? {Significantly lower, Somewhat lower, Little/No change, Somewhat higher, Significantly higher}
We then code responses for Employment, Capital Spending, and Remote Work as 0 if the CFO stated the level was lower, 1 if there was little/no change, and 2 if the level was higher. We back out effects on the the ratio of physical capital and labor using CFO responses about capital spending and labor. If the firm's new level of capital spending was lower (higher) than that of labor, then we say that Physical Capital/Labor has decreased (increased). Similarly, if the new levels of capital spending and labor are the same, then there was no change to Physical Capital/Labor. That is,

$$\text{Physical Capital/Labor} = \begin{cases} 0 \text{ if Capital Spending response} < \text{Employment response} \\ 1 \text{ if Capital Spending response} = \text{Employment response} \\ 2 \text{ if Capital Spending response} > \text{Employment response} \end{cases}$$

Revenue, Employment and Remote Work refer to the level of the variable. Capital Spending refers to "willingness to spend on structures and equipment." The dependent variable is the CFO's response concerning Employment in columns (1) and (2), Capital Spending in columns (3) and (4), and Remote Work in columns (7) and (8). The dependent variable is the Physical Capital/Labor variable in columns (5) and (6). As the dependent variable in each specification has three categories, each column presents results from an ordered logit regression, and coefficients displayed are odds ratios (an odds ratio less (greater) than one indicates a decrease (increase)). Workplace Flexibility, Investment Flexibility, Customer Interactions and Log # Employees (at the end of 2019) are standardized to unit variance. Thus, the odds ratios display the proportional change in the odds of observing a higher response from a standard deviation change in the relevant variable. Detailed variable definitions are in Appendix C.3. Standard errors are clustered at the two-digit NAICS level and displayed in parentheses below the coefficient. ***, **, * denote significance at 1%, 5%, 10%.

Table C.11: Proxies for Financial Flexibility and Customer Interactions in Compustat

	(1)	(2)	(3)	(4)	(5)	(6)
	Employment Growth			Capital Spending Growth		
Customer Interactions	-0.297***	-0.139	-0.222*	-0.277	0.057	0.010
	(0.087)	(0.103)	(0.123)	(0.339)	(0.387)	(0.418)
Lagged Cash/Assets	-0.021	0.184***	0.009	0.191	0.414***	0.318
	(0.128)	(0.020)	(0.134)	(0.465)	(0.084)	(0.471)
Lagged Leverage	-0.072**	0.043	0.015	-0.069	0.280	0.263
	(0.032)	(0.106)	(0.110)	(0.056)	(0.303)	(0.302)
Customer Interactions × Lagged Cash/Assets	0.460		0.394	0.498		0.217
	(0.287)		(0.300)	(1.024)		(1.039)
Customer Interactions × Lagged Leverage		-0.252	-0.192		-0.770	-0.732
		(0.214)	(0.222)		(0.645)	(0.645)
Observations	4,689	4,689	4,689	4,212	4,212	4,212
R-squared	0.096	0.096	0.096	0.057	0.057	0.057
Flexibility Variables & Controls	Yes	Yes	Yes	Yes	Yes	Yes
NAICS-2 FE	Yes	Yes	Yes	Yes	Yes	Yes
State FE	Yes	Yes	Yes	Yes	Yes	Yes

This table examines the interactive role of proxies for financial flexibility and customer interactions in determining employment and capital spending outcomes for Compustat firms for the year 2020. Columns (1) to (3) examines the effect of customer interactions (demand) on employment growth outcomes, conditional on proxies for firms' financial flexibility. In column (1) we interact our demand variable with the firm's lagged cash/assets (lagged CHE/AT in Compustat). In column (2), we interact demand with the firm's lagged leverage (lagged (DLC + DLTT)/AT in Compustat). In column (3), we include both interactive terms. Columns (4) to (6) display the analogous specifications for capital spending growth outcomes. In all columns, we require that a firm have positive assets, non-negative debt, non-missing data for lagged leverage and cash/assets, non-missing employment or capital spending data from fiscal years 2019 and 2020, and a non-missing four-digit NAICS code. In columns (1) to (3), the dependent variable is the log change in employment from 2019 to 2020. In columns (4) to (6), the dependent variable is the log change in capital spending from 2019 to 2020. Each column includes the variables Workplace Flexibility, Investment Flexibility and Log # Employees (at the end of 2019), and both state and NAICS-2 fixed effects. Standard errors are clustered at the four-digit NAICS level and displayed in parentheses below the coefficient. ***, **, * denote significance at 1%, 5%, 10%.

C.2 Model Solutions

The firm's maximization problem is:

$$\max_{L_1,\Delta,L_2,K_2} B[\phi A L_1^{\alpha} + \phi A(K_1 - \Delta)^{\beta}] + A L_2^{\alpha} + A(K_2 + \Delta)^{\beta} - wL_1 - wL_2 - \xi\Delta^2 - K_1 - K_2,$$

$$s.t.\ \ wL_1 + K_1 - \Delta = C.$$

Let $F(L_1, \Delta, L_2, K_2)$ denote the objective function and define $G = wL_1 + K_1 - \Delta - C$. The partial derivatives of F are:

$$\frac{\partial F}{\partial L_1} = \alpha\phi AB{L_1}^{\alpha-1} - w,$$

$$\frac{\partial F}{\partial L_2} = \alpha A{L_2}^{\alpha-1} - w,$$

$$\frac{\partial F}{\partial \Delta} = -\beta\phi AB(K_1 - \Delta)^{\beta-1} + \beta A(K_2 + \Delta)^{\beta-1} - 2\xi\Delta,$$

$$\frac{\partial F}{\partial K_2} = \beta A(K_2 + \Delta)^{\beta-1} - 1.$$

Financial Constraint Not Binding

Setting all partial derivatives of F to 0 gives the unconstrained first-order conditions (FOCs):

$$\frac{\partial F}{\partial L_1} = \alpha\phi AB{L_1}^{\alpha-1} - w = 0,$$

$$\frac{\partial F}{\partial \Delta} = -\beta\phi AB(K_1 - \Delta)^{\beta-1} + \beta A(K_2 + \Delta)^{\beta-1} - 2\xi\Delta = 0.$$

Impact of Workplace Flexibility. We take derivatives with respect to ϕ on both sides of the FOCs and obtain:

$$\frac{\partial L_1}{\partial \phi} = \frac{L_1}{\phi(1-\alpha)} > 0,$$

$$\frac{\partial \Delta}{\partial \phi} = \frac{\beta AB(K_1 - \Delta)^{\beta-1}}{\beta(\beta-1)\phi AB(K_1 - \Delta)^{\beta-2} + \beta(\beta-1)A(K_2 + \Delta)^{\beta-2} - 2\xi} < 0.$$

Impact of Investment Flexibility. We take derivatives with respect to ξ on both sides of the FOCs and obtain:

$$\frac{\partial L_1}{\partial \xi} = 0,$$

$$\frac{\partial \Delta}{\partial \xi} = \frac{2\Delta}{\beta(\beta-1)\phi AB(K_1-\Delta)^{\beta-2} + \beta(\beta-1)A(K_2+\Delta)^{\beta-2} - 2\xi} \lesseqgtr 0 \ if \ \Delta \gtreqless 0.$$

Impact of Demand Shifter. We take derivatives with respect to B on both sides of the FOCs and obtain:

$$\frac{\partial L_1}{\partial B} = \frac{L_1}{\phi(1-\alpha)} > 0,$$

$$\frac{\partial \Delta}{\partial B} = \frac{\beta\phi A(K_1-\Delta)^{\beta-1}}{\beta(\beta-1)\phi AB(K_1-\Delta)^{\beta-2} + \beta(\beta-1)A(K_2+\Delta)^{\beta-2} - 2\xi} < 0.$$

Financial Constraint Binding

The partial derivatives of G are:

$$\frac{\partial G}{\partial L_1} = w,$$

$$\frac{\partial G}{\partial L_2} = 0,$$

$$\frac{\partial G}{\partial \Delta} = -1,$$

$$\frac{\partial G}{\partial K_2} = 0.$$

Let $Y = F - \lambda G$ denote the Lagrange function. The FOCs are:

$$\frac{\partial Y}{\partial \lambda} = wL_1 + K_1 - \Delta - C = 0,$$

$$\frac{\partial Y}{\partial L_1} = \alpha\phi AB{L_1}^{\alpha-1} - (\lambda+1)w = 0,$$

$$\frac{\partial Y}{\partial \Delta} = -\beta\phi AB(K_1-\Delta)^{\beta-1} + \beta A(K_2+\Delta)^{\beta-1} - 2\xi\Delta + \lambda = 0.$$

Impact of Financial Flexibility. We can take derivatives on both sides of the FOCs with respect to C and obtain:

$$\frac{\partial L_1}{\partial C} = \frac{\beta(\beta-1)\phi AB(K_1-\Delta)^{\beta-2} + \beta(\beta-1)A(K_2+\Delta)^{\beta-2} - 2\xi}{\beta(\beta-1)\phi AB(K_1-\Delta)^{\beta-2}w + \beta(\beta-1)A(K_2+\Delta)^{\beta-2}w - 2\xi w + \frac{1}{w}\alpha(\alpha-1)\phi ABL_1^{\alpha-2}} > 0,$$

$$\frac{\partial \Delta}{\partial C} = \frac{\frac{1}{w^2}\alpha(1-\alpha)\phi ABL_1^{\alpha-2}}{\beta(\beta-1)\phi AB(K_1-\Delta)^{\beta-2} + \beta(\beta-1)A(K_2+\Delta)^{\beta-2} - 2\xi + \frac{1}{w^2}\alpha(\alpha-1)\phi ABL_1^{\alpha-2}} < 0.$$

Impact of Workplace Flexibility. We take derivatives on both sides of the FOCs with respect to ϕ and obtain:

$$\frac{\partial L_1}{\partial \phi} = \frac{1}{w}\frac{\partial \Delta}{\partial \phi},$$

$$\frac{\partial \Delta}{\partial \phi} = \frac{\beta AB(K_1-\Delta)^{\beta-1} - \frac{1}{w}\alpha ABL_1^{\alpha-1}}{\beta(\beta-1)\phi AB(K_1-\Delta)^{\beta-2} + \beta(\beta-1)A(K_2+\Delta)^{\beta-2} - 2\xi + \frac{1}{w^2}\alpha(\alpha-1)\phi ABL_1^{\alpha-2}}.$$

The denominator of $\frac{\partial \Delta}{\partial \phi}$ is always negative (as long as $K_1 - \Delta \geq 0$). The sign of the numerator can be either positive or negative. The intuition is that when the financial constraints are binding, labor competes with capital for financial resources. When ϕ is higher, if sufficiently more resources are given to investment, employment could decline, and vice versa.

Impact of Investment Flexibility. We take derivatives on both sides of the FOCs with respect to ξ and obtain:

$$\frac{\partial L_1}{\partial \xi} = \frac{1}{w}\frac{\partial \Delta}{\partial \xi},$$

$$\frac{\partial \Delta}{\partial \xi} = \frac{2\Delta}{\beta(\beta-1)\phi AB(K_1-\Delta)^{\beta-2} + \beta(\beta-1)A(K_2+\Delta)^{\beta-2} - 2\xi + \frac{1}{w^2}A\alpha(\alpha-1)\phi ABL_1^{\alpha-2}}.$$

The signs of $\frac{\partial \Delta}{\partial \xi}$ and $\frac{\partial L_1}{\partial \xi}$ are also the same, and they depend on Δ. Since the denominator of $\frac{\partial \Delta}{\partial \xi}$ is always negative, we have $\frac{\partial \Delta}{\partial \xi} < 0$ if $\Delta > 0$ and $\frac{\partial \Delta}{\partial \xi} > 0$ if $\Delta < 0$.

Impact of the Demand Shifter. We take derivatives on both sides of the FOCs with respect to B and obtain:

$$\frac{\partial L_1}{\partial B} = \frac{1}{w}\frac{\partial \Delta}{\partial B},$$

$$\frac{\partial \Delta}{\partial B} = \frac{\beta\phi A(K_1-\Delta)^{\beta-1} - \frac{1}{w}\alpha\phi AL_1^{\alpha-1}}{\beta(\beta-1)\phi AB(K_1-\Delta)^{\beta-2} + \beta(\beta-1)A(K_2+\Delta)^{\beta-2} - 2\xi + \frac{1}{w^2}\alpha(\alpha-1)\phi ABL_1^{\alpha-2}}.$$

Similar to the case of workplace flexibility, the comparative statics with respect to the demand shifter can be ambiguous when financial constraints are binding. As demand increases, labor and capital compete for financial resources. One of them will increase and the other will decrease; which one will increase depends on the relative marginal returns of capital and labor.

C.3 Data Appendix

This appendix provides detail on the construction of all variables used in the paper. Code and data for our measures of workplace flexibility and investment flexibility, customer interactions and fixed cost share are available at https://github.com/jwb4335/corporate_flexibility. More detail on the construction of other variables is available upon request.

C.3.1 Duke CFO Survey Variables

Revenue/Employment/Capital Spending Forecasts

CFO's forecast of the 12-month ahead percentage change in revenue, employment and capital spending, as answered in the question below.

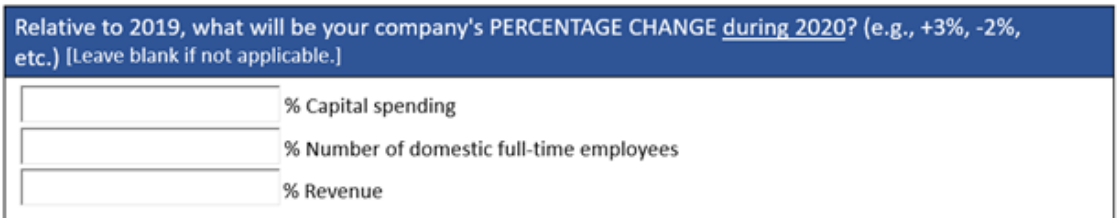

Financial Flexibility

An indicator variable taking a value of one if the CFO answered 2 or above to the question below.

Investment Flexibility

Four-digit NAICS level proxy for a firm's investment flexibility with respect to speed of project completion. We use data from the March 2019 Duke CFO survey to construct a four-digit NAICS code measure of investment ilexibility. Specifically, we define a firm as having flexible investment if they answered "Flexible" or "Very Flexible" to the question below. We then calculate the percentage of firms with investment flexibility at the four-digit NAICS level. Industries with the

highest investment flexibility include beverage, media, apparel stores, and banking, while industries with the lowest investment flexibility include farming, mining, transportation, health care, and wholesale.

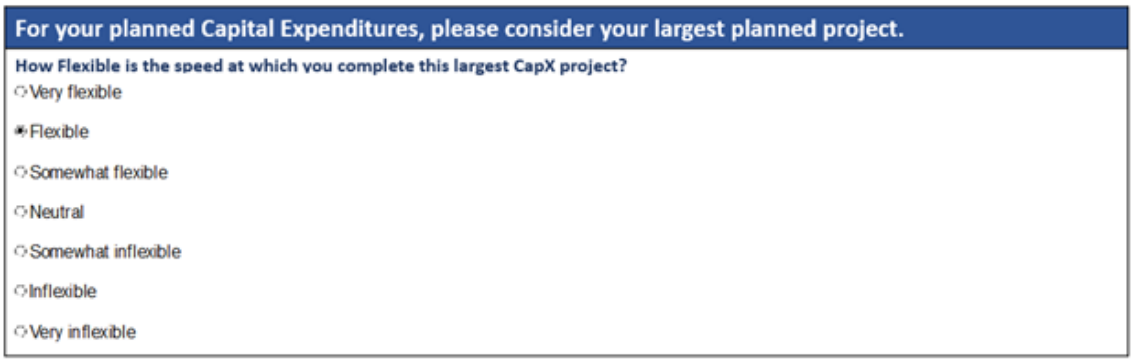

Log # Employees (2019)

The natural logarithm of the firm's number of full-time employees at the end of 2019.

COVID Risk

An indicator variable taking a value of one if the CFO answered with "Medium Coronavirus Risk" or "Large Coronavirus Risk" to the question below.

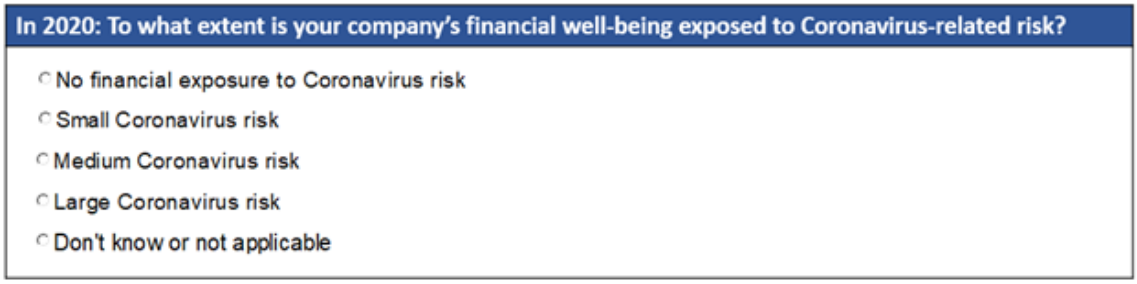

Limited Access to External Capital

An indicator variable taking a value of one if the CFO answer with "Yes, a small amount," "Yes, a moderate amount," or "Yes, a large amount" to the question below.

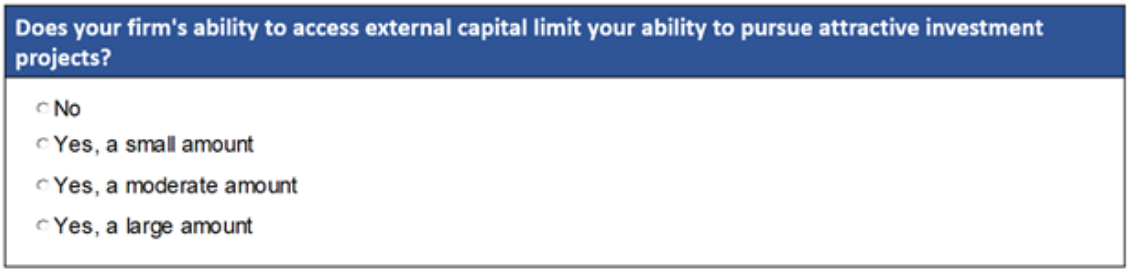

Cash/Assets

Firm's year-end cash to total assets ratio from the March 2020 survey, as answered in the question below.

C.3.2 External Variables

Workplace Flexibility measure from the American Time Use Survey (ATUS)

Four-digit NAICS level proxy for a firm's ability to do work from home. We use data from the 2017-2018 ATUS Leave and Job Flexibilities module (n = 10,040), which asks questions related to workers' ability to perform their job from home. Following Papanikolaou and Schmidt (2022) and Alon et al. (2020), we classify a worker as being able to work from home if they answer yes to these two questions:

- As part of your (main) job, can you work at home?
- Are there days when you work only at home?

Using the Soltas (2019) crosswalk, we aggregate the number of workers that are able to work from home to the four-digit NAICS level. Low workplace flexibility industries include manufacturing and retail; high workplace flexibility industries include professional/scientific services industries.

Workplace Flexibility measure from Dingel and Neiman (2020)

Two-digit NAICS level proxy for a firm's ability to do work from home. This variable is constructed from the O*NET survey and is aggregated from the occupation level to the industry level. Details are available in Dingel and Neiman (2020) and data are available at https://github.com/jdingel/DingelNeiman-workathome.

Customer Interactions

Four-digit NAICS level proxy for the degree of social interactions with customers when they purchase goods and services. We start with the O*NET Work Activities survey which asks workers:

- For your job, how important is performing for people or dealing directly with the public? {1 = Not Important, 2, 3, 4, 5 = Very Important}

This is an occupation-level variable that tracks the importance of *direct* customer interactions in the transactions of goods and services in each industry, similar to those used in Koren and Pető (2020) and Pagano et al. (2021). We convert our measure to a $[0, 1]$ scale. Using the same method as in Dingel and Neiman (2020), we aggregate this variable to the 4-digit NAICS level, using the proportion of occupation-level workers in each industry as weights. We term this variable "Direct Customer Interactions." This measure does not capture customer interactions and the associated impact on customer demand during the COVID-19 health crisis for non-consumer facing industries. For example, airlines have a high Direct Customer Interaction measure, but aircraft manufacturers do not.

We combine our direct measure with the 2012 BEA Input-Output table (the most recent one) to construct a measure for the importance of downstream linkages to industries with a high degree of direct customer interactions. For each industry, we construct downstream output weights to other industries: when industry i is the supplier and industry j is a downstream industry, the weight on industry j is industry i's output supplied to industry j divided by i's total intermediate output. Using these weights, we construct our measure "Indirect Customer Interactions" as the weighted average of downstream industries' Direct Customer Interactions. We combine the direct

and indirect channels into an overall measure for the importance of customer interactions:

$$\text{Customer Interactions}_i = C_i(\text{Direct}_i) + (1 - C_i)(\text{Indirect}_i).$$

The ratios C_i and $1 - C_i$ capture the fraction of an industry's output sold to consumers and to other industries respectively:

$$C_i = \left(\frac{\text{Personal Consumption Expenditures}}{\text{Total Intermediate Output + Personal Consumption Expenditures}} \right)_i,$$

where "Personal Consumption Expenditures" is series F010000 and "Total Intermediate Output" is series T001 in the 2012 Input-Output tables. Airlines are in the top decile of the importance of direct customer interactions (high importance), whereas aircraft manufacturing is in the bottom decile (low importance). Conversely, airlines are slightly above median for the importance of indirect customer interactions, and aircraft manufacturing ranks in the top decile (i.e., aircraft manufacturing sells heavily to customer-facing industries).

Fixed Cost Share

Four-digit NAICS level proxy for the proportion of a firm's operating costs are fixed, as opposed to variable. Using quarterly Compustat data from 1985-2018, and following Anderson et al. (2003) and Chen et al. (2019), we regress log changes in operating costs on log changes in sales at the industry level. The slope coefficient from this regression indicates the fraction of variable costs in total costs (variable cost share). We take our measure of fixed cost share as 1 – variable cost share. Industries with low fixed cost share include those in retail/wholesale, whereas those with high fixed cost shares include mining (for example, oil & gas extraction) and pharmaceutical/medical manufacturing.

Fraction Part-Time

Four-digit NAICS level variable for the percentage of workers in an industry that are part-time. Following Mas and Pallais (2017), we start from the 2010, 2012, 2014, 2016, 2018 supplements of the General Social Survey. Using the Soltas (2019) crosswalk to map from census to NAICS

industry codes, we aggregate to the NAICS-4 level, and take the simple average across years as our final measure. Retail industries have a high fraction of part-time workers (e.g., clothing stores), whereas professional/scientific services have a low fraction of part-time workers (e.g., scientific/R&D services).

Scheduling Autonomy

Four-digit NAICS level proxy for a worker's autonomy in setting their work schedule. Following Mas and Pallais (2017), we use data from the ATUS Leave and Job Flexibilities Module and classify a worker as having scheduling autonomy if they answer yes to the following:

- Do you have flexible work hours that allow you to vary or make changes in the times you begin and end work?

Using the Soltas (2019) crosswalk, we aggregate the number of workers with scheduling autonomy to the four-digit NAICS level. Professional/scientific services have high levels of scheduling autonomy (e.g., architectural/engineering services); whereas manufacturing industries tend to have low levels of scheduling autonomy (e.g., dairy production, among other food production industries).

Unionization

Four-digit NAICS level variable for the percentage of workers in an industry that are union members. Following Hirsch and Macpherson (2003), we start from the 2019 Current Population Survey (CPS) files. Using the Soltas (2019) crosswalk to map from census to NAICS industry codes, we aggregate to the NAICS-4 level. Air transportation industries have high unionization rates, whereas retail (e.g., clothing stores) have low unionization rates.

Scale Inflexibility

Four-digit NAICS level proxy for a firm's inability to easily adjust the scale of production in response to profitability shocks. Using annual Compustat data from 2000-2019, We start from the

firm-level measure described in Gu et al. (2019) and Gu et al. (2021):

$$INFLEX_{i,t} = \frac{\max_{i,t_0,t}\left\{\frac{OPC}{Sales}\right\} - \min_{i,t_0,t}\left\{\frac{OPC}{Sales}\right\}}{\text{stddev}_{i,t_0,t}\left(\Delta \log\left(\frac{Sales}{Assets}\right)\right)},$$

where the numerator is the range of the firm's ratio of operating costs to sales over the previous 20 years, and the denominator is the standard deviation over the previous 20 years of the firms sales to assets ratio (we require firms to have at least 10 years of data to be included). We then take the simple average across firms at the NAICS-4 level as our final measure. Firms in oil & gas extraction have a high degree of inflexibility, whereas firms in plastic and paper product manufacturing have low inflexibility.

Human Coordination

Four-digit NAICS level variable that captures the importance of human interactions in a firm's operations. We use five questions from the O*NET Work Activities and Work Context surveys that asks workers how important the following are to their job: (i) Face-to-face discussions, (ii) Work with group our team, (iii) Contact with others, (iv) Developing and building teams and (v) Communicating with supervisors, peers or subordinates. These questions do not differentiate between remote and in-person activities, so we focus on occupations with above-median importance of physical proximity (dropping below-median occupations).We then aggregate to the four-digit NAICS level using similar methods as in Dingel and Neiman (2020). We define high human coordination industries as those with human coordination importance in the top quartile. Two industries with high Human Coordination are Restaurants and Scheduled air transportation.

Bibliography

Abraham, K. G., Haltiwanger, J. C., Sandusky, K., and Spletzer, J. R. (2018), "Measuring the gig economy: current knowledge and open issues," Working Paper 24950, National Bureau of Economic Research.

Acemoglu, D. and Restrepo, P. (2020), "Robots and jobs: evidence from US labor markets," *Journal of Political Economy*, 128, 2188–2244.

Acemoglu, D., Chernozhukov, V., Werning, I., and Whinston, M. D. (2021), "Optimal Targeted Lockdowns in a Multigroup SIR Model," *American Economic Review: Insights*, 3, 487–502.

Acharya, V. V. and Steffen, S. (2020), "The risk of being a fallen angel and the corporate dash for cash in the midst of COVID," *Review of Corporate Finance Studies*, 9, 430–471.

Acharya, V. V., Almeida, H., Amihud, Y., and Liu, P. (2021), "Efficiency or resiliency? Corporate choice between financial and operational hedging," Working Paper 3792886, Available at SSRN.

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

Aggarwal, R., Dahiya, S., and Yilmaz, U. (2023), "Why do Investors Vote Against Corporate Directors?" *Working Paper*.

Alekseev, G., Amer, S., Gopal, M., Kuchler, T., Schneider, J., Stroebel, J., and Wernerfelt, N. C. (2020), "The effects of COVID-19 on U.S. small businesses: evidence from owners, managers, and employees," Working Paper 27833, National Bureau of Economic Research.

Alexeev, M. and Leitzel, J. (1996), "Rent Shrinking," *Southern Economic Journal*, 62, 90–113.

Almeida, H. and Campello, M. (2007), "Financial constraints, asset tangibility, and corporate investment," *The Review of Financial Studies*, 20, 1429–1460.

Alon, T., Doepke, M., Olmstead-Rumsey, J., and Tertilt, M. (2020), "The impact of COVID-19 on gender equality," Working Paper 26947, National Bureau of Economic Research.

Alvero, A. and Xiao, K. (2023), "Fuzzy bunching," *Working Paper*.

Alvero, A., Ando, S., and Xiao, K. (2023), "Watch what they do, not what they say: Estimating regulatory costs from revealed preferences," *The Review of Financial Studies*, 36, 2224–2273.

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

Anderson, M. C., Banker, R. D., and Janakiraman, S. N. (2003), "Are selling, general, and administrative costs "sticky"?" *Journal of Accounting Research*, 41, 47–63.

Andrei, D. and Carlin, B. I. (2022), "Schumpeterian Competition in a Lucas Economy," Working Paper, Rice University.

Antill, S. (2021), "Are bankruptcy professional fees excessively high?" *Available at SSRN 3554835*.

Appel, I. R., Gormley, T. A., and Keim, D. B. (2016), "Passive investors, not passive owners," *Journal of Financial Economics*, 121, 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, 909–950.

Arnoud, A., Guvenen, F., and Kleineberg, T. (2019), "Benchmarking global optimizers," Tech. rep., National Bureau of Economic Research.

Babenko, I., Boguth, O., and Tserlukevich, Y. (2016), "Idiosyncratic cash flows and systematic risk," *The Journal of Finance*, 71, 425–456.

Babenko, I., Choi, G., and Sen, R. (2019), "Management (of) proposals," *Available at SSRN 3155428*.

Balsam, S., Boone, J., Liu, H., and Yin, J. (2016), "The impact of say-on-pay on executive compensation," *Journal of Accounting and Public Policy*, 35, 162–191.

Balyuk, T., Prabhala, N. R., and Puri, M. (2021), "Small bank financing and funding hesitancy in a crisis: evidence from the Paycheck Protection Program," Working Paper 3717259, Available at SSRN.

Baqaee, D. and Farhi, E. (2020a), "Nonlinear production networks with an application to the COVID-19 crisis," Working Paper 27281, National Bureau of Economic Research.

Baqaee, D. and Farhi, E. (2021), "Supply and demand in disaggregated Keynesian economies with an application to the COVID-19 crisis," *American Economic Review*, (Forthcoming).

Baqaee, D. R. and Farhi, E. (2019), "Replication Data for: Productivity and Misallocation in General Equilibrium," .

Baqaee, D. R. and Farhi, E. (2020b), "Productivity and misallocation in general equilibrium," *The Quarterly Journal of Economics*, 135, 105–163.

Barrero, J. M. (2022), "The micro and macro of managerial beliefs," *Journal of Financial Economics*, 143, 640–667.

Barrero, J. M., Bloom, N., and Davis, S. J. (2020), "COVID-19 Is Also a Reallocation Shock," *Brookings Papers on Economic Activity*, Summer, 329–371.

Barrero, J. M., Bloom, N., and Davis, S. J. (2021), "Why working from home will stick," Working Paper 28731, National Bureau of Economic Research.

Bartik, A. W., Cullen, Z. B., Glaeser, E. L., Luca, M., and Stanton, C. T. (2020), "What jobs are being done at home during the COVID-19 crisis? Evidence from firm-level surveys," Working Paper 27422, National Bureau of Economic Research.

Baye, M. and Hoppe, H. (2003), "The Strategic Equivalence of Rent-Seeking, Innovation, and Patent-Race Games," *Games and Economic Behavior*, 44, 217–226.

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

Ben-David, I., Graham, J. R., and Harvey, C. R. (2013), "Managerial miscalibration," *Quarterly Journal of Economics*, 128, 1547–1584.

Bernardo, A. E., Cai, H., and Luo, J. (2001), "Capital budgeting and compensation with asymmetric information and moral hazard," *Journal of Financial Economics*, 61, 311–344.

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

Bertrand, M. and Schoar, A. (2003), "Managing with style: The effect of managers on firm policies," *The Quarterly journal of economics*, 118, 1169–1208.

Bessembinder, H. and Décaire, P. H. (2021), "Discount rate uncertainty and capital investment," Available at SSRN.

Bick, A., Blandin, A., and Mertens, K. (2021), "Work from home before and after the COVID-19 outbreak," Working Paper DP15000, Center for Economic and Policy Research.

Blonien, P., Crane, A. D., Crotty, K., and De Angelis, D. (2022), "Errors in Shareholder Voting," *Working Paper*.

Bloom, N., Liang, J., Roberts, J., and Ying, Z. J. (2014), "Does Working from Home Work? Evidence from a Chinese Experiment," *Quarterly Journal of Economics*, 130, 165–218.

Bloom, N., Fletcher, R. S., and Yeh, E. (2021), "The impact of COVID-19 on US firms," Working Paper 28314, National Bureau of Economic Research.

Blundell, W., Gowrisankaran, G., and Langer, A. (2020), "Escalation of Scrutiny: The Gains from Dynamic Enforcement of Environmental Regulations," *American Economic Review*, 110, 2558–85.

Bol, J. C. (2011), "The determinants and performance effects of managers' performance evaluation biases," *The Accounting Review*, 86, 1549–1575.

Bonatti, A., Cisternas, G., and Toikka, J. (2017), "Dynamic oligopoly with incomplete information," *The Review of Economic Studies*, 84, 503–546.

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

Brealey, R. and Myers, S. (1996), *Principles of Corporate Finance*, McGrawy Hill Publisher.

Brynjolfsson, E., Horton, J. J., Ozimek, A., Rock, D., Sharma, G., and TuYe, H.-Y. (2020), "COVID-19 and remote work: an early look at US data," Working Paper 27344, National Bureau of Economic Research.

Burgstahler, D. and Dichev, I. (1997), "Earnings management to avoid earnings decreases and losses," *Journal of accounting and economics*, 24, 99–126.

Burtraw, D. (1993), "Bargaining with Noisy Delegation," *RAND Journal of Economics*, 24, 40–57.

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

Campello, M., Graham, J. R., and Harvey, C. R. (2010), "The real effects of financial constraints: evidence from a financial crisis," *Journal of Financial Economics*, 97, 470–487.

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

Charoenwong, B., Kimura, Y., Kwan, A., and Tan, E. (2021), "Capital budgeting, uncertainty, and misallocation," *Working Paper*.

Chen, Q. and Jiang, W. (2004), "Positive hurdle rates without asymmetric information," *Finance Research Letters*, 1, 106–112.

Chen, Z., Harford, J., and Kamara, A. (2019), "Operating leverage, profitability, and capital structure," *Journal of Financial and Quantitative Analysis*, 54, 369–392.

Chetty, R., Friedman, J. N., Olsen, T., and Pistaferri, L. (2011), "Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from Danish tax records," *The quarterly journal of economics*, 126, 749–804.

Chetty, R., Friedman, J. N., Hendren, N., Stepner, M., and The Opportunity Insights Team (2020), "The economic impacts of COVID-19: evidence from a new public database built using private sector data," Working Paper 27431, National Bureau of Economic Research.

Chodorow-Reich, G. (2014), "The employment effects of credit market disruptions: firm-level evidence from the 2008–9 financial crisis," *Quarterly Journal of Economics*, 129, 1–59.

Chowdhury, S. M. and Sheremeta, R. M. (2011), "A generalized Tullock contest," *Public Choice*, 147, 413–420.

Chung, T.-Y. (1996), "Rent-seeking contest when the prize increases with aggregate efforts," *Public Choice*, 87, 55–66.

Clayton, M. C., Hartzell, J. C., and Rosenberg, J. (2005), "The impact of CEO turnover on equity volatility," *The Journal of Business*, 78, 1779–1808.

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

Cooper, R. W. and Haltiwanger, J. C. (2006), "On the nature of capital adjustment costs," *Review of Economic Studies*, 73, 611–633.

Corrado, C., Hulten, C., and Sichel, D. (2009), "Intangible capital and US economic growth," *Review of Income and Wealth*, 55, 661–685.

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, 500–520.

Crawford, V. and Varian, H. (1979), "Distortions of Preferences and the Nash Theory of Bargaining." *Economics Letters*, 3, 203–206.

Crouzet, N. and Eberly, J. C. (2019), "Understanding weak capital investment: the role of market concentration and intangibles," Working Paper 25869, National Bureau of Economic Research.

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

Cziraki, P. and Jenter, D. (2022), "The market for CEOs," *Working Paper*.

Dalton, D. R. and Dalton, C. M. (2008), "Corporate governance in the post Sarbanes–Oxley period: Compensation disclosure and analysis (CD&A)," *Business Horizons*, 51, 85–92.

D'Aspremont, C. and Jacquemin, A. (1988), "Cooperative and Noncooperative R&D in Duopoly with Spillovers," *American Economic Review*, 78, 1133–1137.

De Loecker, J., Eeckhout, J., and Unger, G. (2020), "The rise of market power and the macroeconomic implications," *The Quarterly Journal of Economics*, 135, 561–644.

Décaire, P. (2021), "Capital budgeting and idiosyncratic risk," Working Paper.

Décaire, P. and Sosyura, D. (2022), "On a Spending Spree: The Real Effects of Heuristics in Managerial Budgets," Working Paper.

Dessaint, O., Olivier, J., Otto, C. A., and Thesmar, D. (2021), "CAPM-based company (mis) valuations," *The Review of Financial Studies*, 34, 1–66.

Dessaint, O., Foucault, T., and Frésard, L. (2022), "The Horizon of Investors' Information and Corporate Investment," Working Paper.

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

Ding, W., Levine, R., Lin, C., and Xie, W. (2021), "Corporate immunity to the COVID-19 pandemic," *Journal of Financial Economics*, 141, 802–830.

Dingel, J. I. and Neiman, B. (2020), "How many jobs can be done at home?" *Journal of Public Economics*, 189, 104235.

Duan, Y. and Jiao, Y. (2016), "The role of mutual funds in corporate governance: Evidence from mutual funds' proxy voting and trading behavior," *Journal of Financial and Quantitative Analysis*, 51, 489–513.

Eberly, J. C., Haskel, J., and Mizen, P. (2021), ""Potential Capital", working from home, and economic resilience," Working Paper 29431, National Bureau of Economic Research.

Edmans, A., Gosling, T., and Jenter, D. (2021), "CEO compensation: Evidence from the field," *CEPR Discussion Paper No. DP16315*.

Eichenbaum, M. S., Rebelo, S., and Trabandt, M. (2021), "The macroeconomics of epidemics," *Review of Financial Studies*, 34, 5149–5187.

Eisfeldt, A. L. and Papanikolaou, D. (2013), "Organization capital and the cross-section of expected returns," *Journal of Finance*, 68, 1365–1406.

Ewens, M., Xiao, K., and Xu, T. (2023), "Regulatory costs of being public: Evidence from bunching estimation," *Working Paper*.

Fahlenbrach, R., Rageth, K., and Stulz, R. M. (2021), "How valuable is financial flexibility when revenue stops? Evidence from the COVID-19 crisis," *Review of Financial Studies*, 34, 5474–5521.

Faulkender, M. and Petersen, M. A. (2005), "Does the Source of Capital Affect Capital Structure?" *The Review of Financial Studies*, 19, 45–79.

Favilukis, J. Y., Lin, X., Sharifkhani, A., and Zhao, X. (2020), "Labor force telework flexibility and asset prices: evidence from the COVID-19 pandemic," Working Paper 3693239, Available at SSRN.

Fazzari, S. M., Hubbard, R. G., and Petersen, B. C. (1988), "Financing Constraints and Corporate Investment," *Brookings Papers on Economic Activity*, 1988, 141–206.

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

Fershtman, C., Judd, K., and Kalai, E. (1991), "Observable Contracts: Strategic Delegation and Cooperation," *International Economic Review*, 32, 551–559.

Fink, J., Fink, K. E., Grullon, G., and Weston, J. P. (2010), "What drove the increase in idiosyncratic volatility during the internet boom?" *Journal of Financial and Quantitative Analysis*, 45, 1253–1278.

Fisch, J., Palia, D., and Solomon, S. D. (2018), "Is say on pay all about pay: The impact of firm performance," *Harv. Bus. L. Rev.*, 8, 101.

Fischer, P. E., Gramlich, J. D., Miller, B. P., and White, H. D. (2009), "Investor perceptions of board performance: Evidence from uncontested director elections," *Journal of Accounting and Economics*, 48, 172–189.

Fornaro, L. and Wolf, M. (2020), "COVID-19 coronavirus and macroeconomic policy," *Working Paper*.

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

Foster, F. D. and Viswanathan, S. (1996), "Strategic trading when agents forecast the forecasts of others," *The Journal of Finance*, 51, 1437–1478.

Gabaix, X. and Landier, A. (2008), "Why has CEO pay increased so much?" *The quarterly journal of economics*, 123, 49–100.

Giglio, S., Maggiori, M., Stroebel, J., and Utkus, S. (2021), "The joint dynamics of investor beliefs and trading during the COVID-19 crash," *Proceedings of the National Academy of Sciences*, 118.

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

Gofman, M., Segal, G., and Wu, Y. (2020), "Production networks and stock returns: The role of vertical creative destruction," *The Review of Financial Studies*, 33, 5856–5905.

Gompers, P. A., Kaplan, S. N., and Mukharlyamov, V. (2020), "Private equity and COVID-19," Working Paper 27889, National Bureau of Economic Research.

Gormsen, N. J. and Huber, K. (2020), "Equity Factors and Firms' Perceived Cost of Capital," Available at SSRN.

Gormsen, N. J. and Huber, K. (2022), "Corporate Discount Rates," Available at SSRN.

Graham, J. R. (2022), "Presidential Address: Corporate Finance and Reality," *The Journal of Finance*, 77, 1975–2049.

Graham, J. R. and Harvey, C. R. (2001a), "The theory and practice of corporate finance: Evidence from the field," *Journal of financial economics*, 60, 187–243.

Graham, J. R. and Harvey, C. R. (2001b), "The theory and practice of corporate finance: evidence from the field," *Journal of Financial Economics*, 60, 187–243.

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

Granja, J., Makridis, C., Yannelis, C., and Zwick, E. (2020), "Did the Paycheck Protection Program hit the target?" Working Paper 27095, National Bureau of Economic Research.

Gu, L., Hackbarth, D., and Li, T. (2019), "Inflexibility and leverage," Working Paper 3296926, Available at SSRN.

Gu, L., Hackbarth, D., and Li, T. (2021), "Inflexibility and corporate cash policies," *Working Paper*.

Guerrieri, V., Lorenzoni, G., Straub, L., and Werning, I. (2021), "Macroeconomic Implications of COVID-19: can Negative Supply Shocks Cause Demand Shortages?" *American Economic Review*, (Forthcoming).

Harris, M. and Raviv, A. (1996), "The capital budgeting process: Incentives and information," *The Journal of Finance*, 51, 1139–1174.

Harris, M. and Raviv, A. (1998), "Capital budgeting and delegation," *Journal of Financial Economics*, 50, 259–289.

Harris, M., Kriebel, C. H., and Raviv, A. (1982), "Asymmetric information, incentives and intrafirm resource allocation," *Management science*, 28, 604–620.

Hartzell, J. C. and Starks, L. T. (2003), "Institutional investors and executive compensation," *The journal of finance*, 58, 2351–2374.

Haskel, J. and Westlake, S. (2018), *Capitalism without Capital: The Rise of the Intangible Economy*, Princeton University Press.

Hassan, T. A., Hollander, S., van Lent, L., Schwedeler, M., and Tahoun, A. (2021), "Firm-level exposure to epidemic diseases: COVID-19, SARS, and H1N1," Working Paper 26971, National Bureau of Economic Research.

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, 91–131.

Hensvik, L., Le Barbanchon, T., and Rathelot, R. (2020), "Which jobs are done from home? Evidence from the American Time Use Survey," Working Paper 3574551, Available at SSRN.

Hirsch, B. T. and Macpherson, D. A. (2003), "Union Membership and Coverage Database from the Current Population Survey: Note," *Industrial and Labor Relations Review*, 56, 349–354.

Hirschman, A. O. (1970), *Exit, voice, and loyalty: Responses to decline in firms, organizations, and states*, vol. 25, Harvard university press.

Hirshleifer, J. (1989), "Conflict and Rent-Seeking Success Functions: Ratio vs. Difference Models of Relative Success," *Public Choice*, 41, 101–112.

Holland, K., Lim, C., and Yi, I. (2023), "Shareholder Meeting Uncertainty: Evidence from the Options Market," *Working Paper*.

Hong, H., Kubik, J. D., Wang, N., Xu, X., and Yang, J. (2021), "Pandemics, vaccines and an earnings damage function," Working Paper 27829, National Bureau of Economic Research.

Iliev, P., Lins, K. V., Miller, D. P., and Roth, L. (2015), "Shareholder voting and corporate governance around the world," *The Review of Financial Studies*, 28, 2167–2202.

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

Jagannathan, R., Matsa, D. A., Meier, I., and Tarhan, V. (2016), "Why do firms use high discount rates?" *Journal of Financial Economics*, 120, 445–463.

Jones, S. (1989), "Have Your Lawyer Call My Lawyer," *Journal of Economic Behavior and Organization*, 11, 159–174.

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

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

Kim, H. and Kung, H. (2017), "The asset redeployability channel: How uncertainty affects corporate investment," *The Review of Financial Studies*, 30, 245–280.

Kogan, L., Papanikolaou, D., Schmidt, L. D., and Seegmiller, B. (2021), "Technology-skill complementarity and Labor displacement: evidence from linking two centuries of patents with occupations," *Working Paper*.

Koren, M. and Pető, R. (2020), "Business disruptions from social distancing," *PLOS ONE*, 15, 1–14.

Krüger, P., Landier, A., and Thesmar, D. (2015), "The WACC fallacy: The real effects of using a unique discount rate," *The Journal of Finance*, 70, 1253–1285.

Lamont, O. A. (2000), "Investment plans and stock returns," *Journal of Finance*, 55, 2719–2745.

Landier, A. and Thesmar, D. (2020), "Earnings expectations during the COVID-19 crisis," *Review of Asset Pricing Studies*, 10, 598–617.

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

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

Liu, T. (2021), "Bargaining with private equity: implications for hospital prices and patient welfare," *Available at SSRN 3896410*.

Lyman, N. (2023), "Reputation, Entrenchment, and Dynamic Managerial Incentives," *Working Paper*.

Mas, A. and Pallais, A. (2017), "Valuing Alternative Work Arrangements," *American Economic Review*, 107, 3722–59.

Mas, A. and Pallais, A. (2020), "Alternative Work Arrangements," *Annual Review of Economics*, 12, 631–658.

Matveyev, E. (2017), "Sorting in the US corporate executive labor market," *Working Paper*.

McCrary, J. (2008), "Manipulation of the running variable in the regression discontinuity design: A density test," *Journal of econometrics*, 142, 698–714.

McDonald, R. L. (2000), "Real Options and Rules of Thumb in Capital Budgeting," in *Project Flexibility, Agency, and Competition: New Developments in the Theory and Application of Real Options*, pp. 13–33, Oxford University Press.

McFadden, D. (1989), "A method of simulated moments for estimation of discrete response models without numerical integration," *Econometrica: Journal of the Econometric Society*, pp. 995–1026.

Mehran, H. (1995), "Executive compensation structure, ownership, and firm performance," *Journal of financial economics*, 38, 163–184.

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

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*, vol. 120, p. 80, HeinOnline.

Olley, G. S. and Pakes, A. (1996), "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64, 1263–1297.

Pagano, M., Wagner, C., and Zechner, J. (2021), "Disaster resilience and asset prices," *Working Paper*.

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

Papanikolaou, D. and Schmidt, L. D. W. (2022), "Working remotely and the supply-side smpact of COVID-19," *The Review of Asset Pricing Studies*, 12, 53–111.

Pastor, L. and Veronesi, P. (2009), "Learning in financial markets," *Annu. Rev. Financ. Econ.*, 1, 361–381.

Patatoukas, P. N. (2012), "Customer-base concentration: Implications for firm performance and capital markets," *The accounting review*, 87, 363–392.

Pinnington, J. (2022), "Are Passive Investors Biased Voters?" *Working Paper*.

Poterba, J. M. and Summers, L. H. (1995), "A CEO survey of US companies' time horizons and hurdle rates," *MIT Sloan Management Review*, 37, 43.

Ramelli, S. and Wagner, A. F. (2020), "Feverish Stock Price Reactions to COVID-19," *Review of Corporate Finance Studies*, 9, 622–655.

Rubinstein, A. (1982), "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 50, 97–109.

Rubinstein, A. and Wolinsky, A. (1985), "Equilibrium in a market with sequential bargaining." *Econometrica*, 53, 1133–1150.

Saez, E. (2010), "Do taxpayers bunch at kink points?" *American economic journal: Economic policy*, 2, 180–212.

Schelling, T. (1956), "An Essay on Bargaining," *American Economic Review*, 46, 281–306.

Segendorff, B. (1998), "Delegation and Threat in Bargaining." *Games and Economic Behavior*, 23, 266–283.

Sharpe, S. A. and Suarez, G. A. (2021), "Why isn't business investment more sensitive to interest rates? evidence from surveys," *Management Science*, 67, 720–741.

Shleifer, A. and Vishny, R. W. (1997), "A Survey of Corporate Governance," *The Journal of Finance*, 52, 737–783.

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

Smith, A. A. (2016), *Indirect Inference*, pp. 1–6, Palgrave Macmillan UK, London.

Sobel, J. (1981), "Distortion of Utilities and the Bargaining Problem," *Econometrica*, 49, 597–619.

Soltas, E. (2019), "Census–NAICS 2012 Industry Code Crosswalk," https://doi.org/10.7910/DVN/O7JLIC.

Taylor, L. A. (2010), "Why Are CEOs Rarely Fired? Evidence from Structural Estimation," *The Journal of Finance*, 65, 2051–2087.

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

Tullock, G. (1980), "Efficient rent seeking," in *Toward a theory of the rent-seeking society*, eds. J. M. Buchanan, R. D. Tollison, and G. Tullock, Texas A&M University Press: College Station.

Whited, T. M. and Wu, G. (2006), "Financial constraints risk," *The Review of Financial Studies*, 19, 531–559.

Biography

John W. Barry completed his Leaving Certificate at Coláiste Iognáid ("The Jes"), in Galway, Ireland. He received his B.A. in Philosophy, Politics, Economics and Sociology from Trinity College Dublin. He received his M.S. in International Trade, Finance and Development from the Barcelona School of Economics, Universitat Autònoma de Barcelona and Universitat Pompeu Fabra. He received his Ph.D. in Finance, from the Fuqua School of Business at Duke University.