Managerial Beliefs and Agency Conflicts[1]

Dardan Gashi[2]

Swiss Finance Institute, USI Lugano

May, 2026

Abstract

I study the role of managerial beliefs for executive compensation and corporate investment. I develop a dynamic investment model with agency conflicts in which the manager can be non-rational. The manager can be optimistic/pessimistic and overconfident/underconfident about future cash flows. Under the optimal contract, managerial beliefs affect incentive provision, firm value and investment policy. Using microdata on managerial sales forecasts, I structurally estimate the model at a granular level and document substantial heterogeneity in managerial biases. Consistent with the model, the contracts observed in the data optimally accommodate managerial beliefs and improve shareholder value. Replacing current managers with rational ones would have only a marginal impact. These results help rationalize the persistence of behavioral biases among top executives in equilibrium.

JEL Codes: D84, D86, E22, J33.

[1]I am deeply grateful to Laurent Frésard, Loriano Mancini, Lukas Schmid and Enrique Schroth for the continuous support and guidance. I also thank Anthony Cookson, Ricardo Delao, Jason Roderick Donaldson, Patrick Gagliardini, William Grieser, Sebastian Gryglewicz, David Hirshleifer, Gerard Hoberg, Wenhao Li, Mete Kilic, Arthur Korteweg, Antonio Mele, Fausto Panunzi, Giorgia Piacentino, Alejandro Rivera, Roberto Steri, Lucian Taylor, Tiange Ye, and seminar participants at SFI USI Lugano and USC Marshall for their helpful comments. I am thankful to John Graham for his support in accessing the Duke CFO Business Outlook Survey. I gratefully acknowledge the financial support from the Swiss Finance Institute and the Swiss National Science Foundation under the SNF grant 100018_204561.

[2]Dardan Gashi is with the Swiss Finance Institute and the Institute of Finance, USI Lugano, Via Giuseppe Buffi 13, CH-6900 Lugano, Switzerland. E-mail: dardan.gashi@usi.ch.

1 Introduction

Economic agents display behavioral biases, and much effort is spent understanding and mitigating their implications. Top executives are no exception, and their characteristics and beliefs explain variation in corporate practices (Malmendier and Tate, 2015, Graham et al., 2011, Graham et al., 2013). In modern corporations, boards design executive compensation to align interests between shareholders and managers. Despite the public debate on the level of executive compensation for incentive provision, little is known about *how* compensation contracts should accommodate non-rational managers, and *whether* the observed contracts do so. How do managerial subjective beliefs affect executive compensation? What are the implications for firm value and investment?

While subjective expectations have been documented across a wide range of economic settings, disagreement in the principal-agent problem remains less well understood. Incentive schemes must be tailored to managers' own beliefs about the firm's cash flow to induce effort provision.[1] Alternatively, even a value maximizing biased manager will take suboptimal decisions and destroy shareholders' value. Yet, financial theory has made limited progress in developing tractable models characterizing the optimal contract and its implications when the principal and the agent disagree about the future outcomes of the firm. Empirically, the lack of data on managers' subjective expectations has largely precluded a quantitative assessment of how belief distortions jointly shape executive compensation, firm value, and

[1]Board members and compensation consultants regularly discuss their work in crafting and tailoring incentive structures for top management through Proxy Statements (SEC DEF 14A) and research reports.

corporate policies.

To guide my work, I start with a dynamic moral hazard model with capital accumulation in which an investor (the principal) hires a manager (the agent) to operate a firm. In this model, agency conflicts arises because the manager can take hidden costly action (effort choices) that affect earnings. I extend this standard framework by introducing disagreement between the principal and the agent. I allow the manager to hold biased beliefs about the firm's future cash flows. The manager may be biased about both the first moment (optimism or pessimism) and the second moment (over- or underconfidence) of cash flows. Two latent characteristics that are well documented in the psychology literature and have received much attention in economics and finance in the recent past. An optimistic (pessimistic) manager overestimates (underestimates) expected cash flows, while an overconfident (underconfident) manager underestimates (overestimates) the volatility. These distortions directly affect incentive compatibility by affecting the manager's perceived value of promised utility, which in turn governs the strength of incentives required to induce effort. The optimal contract specifies, as a function of the history of firm's performance: (i) the agent's compensation, (ii) whether the contract is terminated, and (iii) the level of investment in the firm. Termination involves liquidation of the firm or firing the agent. As usual, the agent is paid only after sufficient good performance and one of the state variables in the model is its "continuation utility" in the form of deferred compensation. The biases of the agent directly shape incentive compatibility and therefore, the dynamics of this state variable. Optimistic and overconfident managers, who perceive a higher value of deferred compensation, require weaker incentives to exert effort. This relaxes the firm's incentive constraint and frees resources for productive investment. Conversely, pessimistic or underconfident managers demand stronger incentives, amplify the agency conflict and crowd out resources.

The model yields several key implications. First, managerial beliefs shape both the level and the structure of executive compensation by affecting the incentive compatibility constraint. Deferred compensation is lower for optimistic and overconfident managers, who are easier to motivate, and higher for pessimistic or underconfident managers, who require stronger incentives. Following a sequence of positive shocks, optimistic and overconfident managers are paid earlier, as the likelihood of termination declines; the opposite is true for pessimistic and underconfident managers. Second, shareholders optimally design compensation contracts to accomodate for managerial biases. Optimistic and overconfident managers are less costly to incentivize, while pessimistic and underconfident managers command higher expected pay. These dynamics evolve with the state variable capturing the manager's continuation utility. The benefit from exploiting optimism displays a hump-shaped pattern: it is small when the firm is near termination, rises after positive shocks as firm value improves, and declines again in good states when the manager's ownership stake becomes large. Overall, the model highlights the importance of jointly considering biases in both the mean and variance of beliefs to understand incentive provision under moral hazard. The first-best allocation arises more easily when the manager is easier to incentivize, underscoring how belief distortions can relax or tighten the agency conflict which affect firm value and investment decision.

To confront the model with the data, I construct a novel dataset of managerial subjective expectations by combining guidance data from IBES and FactSet with firm fundamentals and compensation data from Compustat and Execucomp. These sources report managers' own forecasts of future sales, disclosed mainly through earnings calls and press releases. Managerial guidance captures the decision makers' beliefs about the firm's prospects, the object of interest in the theoretical framework. The resulting dataset provides broad coverage of U.S. publicly traded firms over the period 2005–2019. A concern with guidance data is strategic manipulation: managers may distort long-horizon forecasts to influence market perceptions, meet or beat expectations. To mitigate this concern, I restrict attention to short-term forecasts issued within twelve months of the fiscal year-end. Short-horizon disciplines reporting behavior, as managers face immediate consequences if realizations deviate substantially from guidance. This choice finds support in Ma et al., 2024, who document that short-term managerial guidance exhibits statistical properties comparable to survey-based expectations from the Bank of Italy, where strategic disclosure motives are absent. As a robustness check, I replicate the main analysis using the Duke CFO Business Outlook Survey, conducted in partnership with the Federal Reserve Banks of Richmond and Atlanta. This survey directly elicits executives' beliefs about their own firms' prospects through questionnaires. The trade-off is a substantially smaller sample, as confidential firm identifiers are available only for a subset of respondents. Nevertheless, the consistency of results across these two independent data sources strengthens confidence that the estimated belief distortions reflect genuine managerial expectations.

I start by documenting a set of stylized facts on managerial forecasts, belief distortions, and their relationship to executive compensation. Using the newly constructed dataset, I first establish that managerial sales forecasts are highly informative about realized outcomes. Regressing actual sales on forecasted sales yields slope coefficients close to unity and high R-squared values, indicating that guidance reflects genuine expectations about firm fundamentals rather than noise. This finding validates the empirical content of managerial forecasts as proxies for subjective beliefs. Afterwards, I examine whether these belief distortions map into observable differences in incentive provision. Consistent with the model's predictions, I find that optimistic (pessimistic) managers receive weaker (stronger) pay-for-performance incentives—their compensation packages place less weight on deferred and equity-based components. Conversely, underconfident (overconfident) managers are granted stronger (weaker) incentive schemes. Taken together, these empirical regularities provide direct support for the core mechanisms of the model.

To run a quantification exercise and structurally estimate the model, I derive the dynamics of the model's state variables and discretize them to obtain a tractable empirical representation. The resulting system can be expressed in state-space form, estimated by maximum likelihood using an Unscented Kalman Filter (UKF). This approach allows the model to be estimated at a granular level across homogeneous groups of firms, yielding rich cross-sectional variation in the deep structural parameters. I form small groups of firms that operate in the same industry and share similar characteristics in terms of forecast errors and excess forecast-error volatility (Frésard et al., 2023; Gryglewicz et al., 2021; Gashi, 2025).

The estimation results reveal substantial cross-sectional heterogeneity in managerial beliefs across firms and industries. On average, managers exhibit mild pessimism and overconfidence. Quantitatively, the estimated mean of the first-moment belief parameter α is negative, at -0.009, while the average second-moment belief parameter β is also negative, at -0.006. The first-moment bias displays substantial industry heterogeneity, with some sectors characterized by optimism and others by pessimism. By contrast, estimates of β are mostly negative, indicating that managers tend to underestimate the volatility of future outcomes. Taken together, these findings suggest that managerial beliefs differ along both dimensions: managers are, on average, overconfident about uncertainty, while their expectations about mean outcomes vary across industries.

At the industry level, the estimated belief distortions reveal substantial heterogeneity. First-moment biases vary in both sign and magnitude: managers in consumer goods, manufacturing, and the residual category display mildly optimistic beliefs, while healthcare is close to the rational benchmark and technology exhibits systematic pessimism. By contrast, second-moment distortions are more uniform. The estimates of β are negative across industries, indicating that managers tend to underestimate the volatility of future outcomes. This overconfidence is strongest in manufacturing and is also present, though more moderately, in healthcare, technology, and consumer goods.

These patterns have direct implications for incentive provision. In industries where optimism and overconfidence reinforce each other, the incentive compatibility constraint is relaxed and the optimal contract requires lower incentive intensity relative to the rational benchmark. This is the case in consumer goods, manufacturing, and the 'Others' category. Healthcare lies close to the benchmark, as mild pessimism is largely offset by overconfidence. Technology stands out as the main sector in which belief distortions increase contracting costs: managerial pessimism dominates the mitigating effect of overconfidence, raising the incentive intensity required to induce effort. Overall, the evidence supports the model's prediction that first- and second-moment beliefs jointly determine the cost and structure of executive compensation.

I then examine whether the estimated incentive distortions vary systematically with observable firm characteristics. The results indicate that belief-induced contracting costs are concentrated among younger and smaller firms, as well as among firms with greater intangible intensity and tighter financing constraints. By contrast, firms with higher institutional ownership exhibit significantly smaller distortions, consistent with a role for monitoring and governance in disciplining the effects of managerial beliefs. These cross-sectional patterns complement the industry-level evidence and suggest that belief heterogeneity is not randomly distributed across firms, but is instead related to economically meaningful characteristics of the corporate environment.

A natural question is whether the estimated belief parameters are statistically necessary to explain the data, or whether a simpler model with fully rational managers would provide an equally good fit. To address this, I conduct a series of likelihood ratio tests comparing the unrestricted model against three nested specifications: one that restricts optimism to zero ($\alpha = 0$), one that restricts confidence to its rational benchmark ($\beta = 0$), and one that imposes full rationality ($\alpha = \beta = 0$). The results strongly favour the unrestricted model as the alternative models are widely rejected by the data. Most importantly, the joint hypothesis of managerial rationality is rejected in nearly three-quarters of the sample. These findings carry two important implications. First, they establish that a model incorporating belief provides a superior fit to the data relative to the standard rational-expectations benchmark commonly employed in the dynamic contracting literature. The additional parameters capturing beliefs distorsions are not redundant: they absorb systematic variation in firm behavior that would otherwise remain unexplained. Second, the results indicate that managerial biases are economically useful in explaining observed compensation structures. The fact that executive pay responds to belief distortions precisely as the theory predicts suggests that boards design contracts that optimally account for managerial expectations. This is consistent with the view that observed compensation packages reflect the solution to a well-defined contracting problem in which belief plays a central role. The widespread rejection of rationality across industries, combined with the tight alignment between predicted and observed compensation patterns, provides compelling evidence that the model captures the salient features of the data and constitutes an appropriate framework for the counterfactual analysis that follows.

Next, I quantify the role of managerial beliefs. Consistent with the model, the observed compensation packages in the data differ from those implied by a fully rational benchmark. This raises a series of natural questions: if managers are biased, how much value is created under the optimal contract? How costly are these biases? Why do shareholders retain them in equilibrium? Which firms should try to hire executives with different characteristics?
To address these questions, I first consider a counterfactual that speaks directly to the value of designing optimal contracts. I compare the optimal contract, which accommodates managerial beliefs, with a contract designed as if the manager were fully rational. This exercise isolates the value of board sophistication: rather than contracting with a biased manager as if she were rational, a sophisticated board recognizes the manager's beliefs and designs incentives accordingly. Two cases could arise, depending on how managerial beliefs affect the incentive compatibility constraint. In the first case, beliefs make the manager relatively cheap to incentivize. This occurs, for example, when the manager exhibits optimism or overconfidence. A contract designed for a rational agent still induces effort, but the principal fails to exploit the manager's beliefs. In the second case, beliefs make the manager relatively expensive to incentivize. This occurs for example when pessimism or underconfidence arise. A rational-agent contract would then fail to satisfy incentive compatibility, leading to shirk-

ing and resource diversion. In this scenario, accounting for managerial beliefs is crucial to prevent value destruction. Quantitatively, sophisticated boards increase firm value by 13.9% and corporate investment by 11.3% relative to the naive benchmark. These gains are driven primarily by the large losses that arise when naive contracts fail to induce effort.

I then examine whether shareholders would benefit from replacing biased managers with fully rational ones. In this counterfactual, each firm is assumed to be run by a manager with rational beliefs, $\alpha = \beta = 0$. The analysis shows that firm value would increase by only 0.1 percent on average, while investment policies would remain virtually unchanged. This magnitude is considerably smaller than previous estimates (e.g., Barrero, 2022; Ma et al., 2024), because in my model the optimal contract endogenizes the disagreement between shareholders and managers. The modest gain represents a gross benefit that must be weighed against the likely search, turnover, and matching costs of hiring a new manager. Taken together, these results rationalize the persistence of behavioral biases in corporate settings and underscore the second-order nature of their average welfare effects.

While the average effect of belief distortions on firm value is small, this average masks meaningful cross-sectional heterogeneity. At one extreme, some firms benefit from managerial biases because these biases make executives cheaper to incentivize. For these firms, replacing the biased manager with a rational one would reduce firm value by 0.15% and investment by 0.03%. At the other extreme, some firms are harmed by managerial biases because these biases tighten the incentive compatibility constraint. For these firms, replacing the biased manager with a rational one would increase firm value by 0.06% and investment by 0.02%. I further show that these effects are state-dependent. At current average levels of managerial claims, the welfare impact of biases is small. However, as the manager's continuation utility grows, for instance following a sequence of positive performance shocks, distortions in investment policy become substantially larger.

Finally, I assess the robustness of the main findings using also survey data. To address this, I replicate the structural estimation using the Duke CFO Business Outlook Survey, which elicits executives' expectations through confidential questionnaires that are not disclosed to the market. Despite the substantially smaller sample size, the estimated belief parameters exhibit strong qualitative and quantitative agreement across the two independent data sources: both yield positive estimates of optimism and underconfidence. This concordance provides external validation for the main results in the paper.

1.1 Related Literature

This paper contributes to three main strands of literature.

First, it relates to the growing literature on expectations and behavioural corporate finance. Previous studies have examined how firms form expectations about macroeconomic variables such as inflation (Coibion et al., 2018), GDP growth (Tanaka et al., 2020), aggregate stock market returns (Ben-David et al., 2013), and uncertainty (Guiso and Parigi, 1999). I instead study firms' own subjective sales forecasts and how these beliefs shape their corporate decisions. Malmendier and Tate, 2015 reviews this literature. More recently, Ma et al., 2024 analyse distortions in expectations as a source of misallocation, and Barrero, 2022 studies the micro and macro implications of belief heterogeneity. The most closely related papers are Barrero, 2022, Ma et al., 2024, C. A. Otto, 2014, and Gervais et al., 2011. Other relevant contributions include Stein, 1996, Graham et al., 2013, Graham et al., 2015, Ben-David et al., 2013, Bordalo et al., 2017, and Bordalo et al., 2018.[2] I bring a novel dimension to this literature by introducing executive compensation as a largely neglected channel through which shareholders can mitigate the effects of managerial belief distortions. By jointly analysing managerial expectations and optimal contracting, the paper connects the behavioural literature on firm beliefs with the broader debate on the determinants of

[2]Related work in asset pricing has increasingly incorporated subjective beliefs (Greenwood and Hanson, 2015, Greenwood and Shleifer, 2014, Bouchaud et al., 2019, De La O and Myers, 2021, De La O and Myers, 2024).

incentive provision.

Second, the paper contributes to the literature on dynamic contracting. Most studies in this area analyse agency conflicts without allowing for disagreement between the principal and the agent (see DeMarzo and Sannikov, 2006, Biais et al., 2007, Sannikov, 2008, Miao and Rivera, 2016, Malenko, 2019, Biais et al., 2010, Dittmann et al., 2010, DeMarzo et al., 2012, He, 2009, He, 2011). In contrast, a smaller group of papers explicitly study agency conflicts under belief heterogeneity, such as Rivera, 2023, Gao et al., 2017, and Adrian and Westerfield, 2009. My framework is also related to Ai et al., 2023, Ai and Li, 2015, and Bolton et al., 2011, who study optimal investment and compensation under limited commitment. This paper extends this literature by developing a tractable model that characterizes the optimal contract when the principal and the agent disagree, and by analysing how belief distortions shape the incentives, optimal investment policies, and firm value.

Finally, the paper contributes to the empirical literature that estimates dynamic moral hazard models to quantify the role of agency conflicts (Dittmann and Maug, 2007, L. A. Taylor, 2013, Nikolov and Whited, 2014, Gayle and Miller, 2009, Page, 2018, Nikolov et al., 2021). I develop a tractable methodology that bridges the gap between theory and empirics, easing the mapping of continuous-time dynamic contracting models to the data. The estimation framework allows for heterogeneity in behavioural parameters and can be implemented at a granular level, producing the first quantitative evidence on how managerial beliefs affect incentives, firm value, and investment across industries.

The rest of the paper is organized as follows. Section 2 introduces the theoretical setup, and Section 3 presents the solution of the model. Section 4 and 5 describe the data construction and stylized facts. Section 6 outlines the estimation procedure. Section 7 discusses the results, and Section 8 concludes.

2 The Model

This section develops the theoretical framework that links managerial beliefs, executive compensation, and corporate investment. The model builds on a continuous-time dynamic moral hazard setting with capital accumulation, extending the standard contracting framework of DeMarzo and Sannikov, 2006 to allow for biased managerial expectations and capital accumulation. Managers may be optimistic or pessimistic about future cash flows (first-moment bias) and over- or underconfident about their volatility (second-moment bias). These belief distortions modify the perceived continuation value of the contract and therefore alter the incentive compatibility constraint. By embedding these biases into a dynamic contracting environment, the model highlights how subjective expectations affect both the level and composition of pay and, ultimately, investment decisions. The framework also provides a natural structure for the empirical estimation in later sections, as it yields a state-space representation that can be taken to the data.

2.1 Assumptions

Throughout the paper, time is continuous and uncertainty is modeled by a probability space $(\Omega, \mathcal{F}, F, P)$ with the filtration $F = \{\mathcal{F}_t : t \geq 0\}$, satisfying the usual conditions. I consider a standard principal-agent model in which the risk-neutral owner of a firm (the principal) hires a risk-neutral manager (the agent) to operate the firm's assets. The manager is more impatient than the investor. In the model, firm performance depends on investment and effort provision which entails a monetary cost. Agency problems arise because the manager decides the effort provision and can divert part of the resources. Investment decision is instead fully observable and therefore contractable.

The firm employs capital to produce output, whose price is normalized to one. At any time $t \geq 0$, cashflows are proportional to the capital stock K_t. Following He, 2009, DeMarzo et al., 2012, Bolton et al., 2011, I consider that the firm's capital stock $\{K_t\}_{t\geq 0}$ evolves according to

$$dK_t = (i_t - \delta_k)K_t dt \tag{1}$$

where i_t is the investment rate [3] and $\delta_k > 0$ is the depreciation rate.

Firm output dY_t is proportional to capital K_t but uncertain due to the presence of shocks, and over the increment dt follow [4]

$$dY_t = dA_t K_t \tag{2}$$

where the true process dA_t evolves

$$dA_t = \mu e_t A_t dt + \sigma_a A_t dZ_t \tag{3}$$

where $\mu > 0$ is a constant, $\sigma_a > 0$ is a constant volatility parameter, $e_t \in \{0, 1\}$ is the managerial binary effort provision, and $\{Z\} = \{Z\}_{t\geq 0}$ is a standard Brownian motion. The manager believes that instead dA_t follows

$$dA_t = (1+\alpha)\mu e_t A_t dt + (1+\beta)\sigma_a A_t dZ_t^M \tag{4}$$

where α and β are constants. Moreover, $\alpha > 0$ (< 0) captures optimism (pessimism), and $\beta < 0$ (> 0) captures overconfidence (underconfidence). Investment entails costs $\mathcal{I}(i_t, K_t, A_t)$. We assume that the investment cost is homogeneous of degree one in K_t and A_t for tractability. Therefore, the cost to invest at proportional to $A_t K_t$, as in Gryglewicz et al., 2021, Cooper, 2006, Bustamante et al., 2022, Hackbarth and Johnson, 2015, for tractability. That is, we have that $\mathcal{I}(i_t, K_t, A_t) = \mathcal{C}(i_t)A_t K_t$, where we assume $\mathcal{C}$ is increasing and convex in its

[3]For the problem to be well defined, I consider the case where $i_t \in [0; i_{max}]$ with $i_{max} < r + \delta_k - \mu$ where r is the discount rate of the principal.

[4]I can interpret this linear production function as a reduced form setting with costant returns to scale with labour as a variable factor in production. The optimal labour will then be proportional to the capital stock after following the firm static problem. Moreover, in this setup dA_t captures future demand, input wages, and productivity.

arguments. Unless otherwise mentioned and without loss of generality, I consider quadratic costs of investment function

$$\mathcal{C}(i_t) = \tfrac{1}{2}\lambda_k i_t^2 \tag{5}$$

where $\lambda_k > 0$ is a constant and i_{max} is large enough to ensure that investment is interior at all times. The manager is protected by limited liability, does not accept negative payments from the principal dC_t, and cannot be asked to cover the investment cost out of her own pocket. As in DeMarzo and Sannikov, 2006, DeMarzo et al., 2012, Biais et al., 2007, the agent is more impatient than the principal and has a discount rate $\gamma > r$. As a result, the principal cannot indefinitely postpone payments to the agent. The agent possesses an outside option normalized to zero and maximizes the present value of her expected payoffs.[5] Because the agent is protected by limited liability, her continuation value can never fall below her outside option, in which case she would profit from leaving the firm. Her employment starts at time $t = 0$ and is terminated at an endogenous stopping time τ, at which point the firm is liquidated. At the time of liquidation, the principal recovers a fraction $R > 0$ of assets, valued at $RA_\tau K_\tau$. Liquidation is inefficient and generates deadweight losses.

2.2 The Agency Problem

I now consider an agency conflict induced by the separation of ownership and control. The firm's investors hire an agent to operate the firm. The agent's action $e_t \in \{0, 1\}$ determines the expected rate of output per unit of capital, so it affects the process A_t. The agent controls the drift, but not the volatility of the process A. Note that the firm can incur negative shocks and while these losses can accrue at an unbounded rate given the Brownian motion, I will show that the optimal contract with agency bounds the cumulative losses of the firm by optimally invoking termination.

[5] This is to abstract from the participation constraint, which would simply imply an initial bonus payment to encourage the manager to enter the contract.

When the agent takes the action a_t, he enjoys private benefits at the rate $\phi(1-e_t)A_tK_tdt$, where $0 \leq \phi \leq 1$. The action can be interpreted as an effort choice; due to the linearity of private benefits, our framework is also equivalent to the binary effort setup in which the agent can *shirk*, $a = 0$, or *work*, $a = 1$. Alternatively, we can interpret $1 - a_t$ as the fraction of cash flow that the agent diverts for his private benefit, with ϕ equal to the agent's net consumption per dollar diverted. In either case, ϕ represents the severity of the agency problem and, as we show later, captures the minimum level of incentives required to motivate the agent.

Investors have unlimited wealth and are risk-neutral with discount rate $r > 0$. The agent is also risk-neutral, but with a higher discount rate $\gamma > r$. The impatience assumption avoids the scenario in which investors indefinitely postpone payments to the agent. The agent has no initial wealth and has limited liability, so investors cannot pay negative wages to the agent. If the contract is terminated, the agent's reservation value, which is associated with his next-best employment opportunity, is normalized to zero.

I assume that the firm's capital stock, K_t, and its (cumulative) output, Y_t, are observable and contractible. Therefore, investment i_t and A_t are also contractible. To maximize firm value, investors offer a contract that specifies the firm's investment policy i_t, the agent's cumulative compensation C_t, and a termination time τ, all of which depend on the history of the agent's performance, which is given by the productivity process A_t. The agent's limited liability requires the compensation process C_t to be nondecreasing. I let $\Pi = (i_t, C, \tau)$ represent the contract and leave further regularity conditions to Appendix A.

2.2.1 The Managerial Problem

Given the contract Π, the agent chooses an action process $\{e_t \in \{0,1\} : 0 \leq t < \tau\}$ to solve

$$\max_{\{e_t \in \{0,1\}:0\leq t<\tau\}} E^M \left[\int_0^\tau e^{-\gamma t} \left(dC_t + \phi(1-e_t)\mu A_t K_t dt \right) \right] \tag{6}$$

where $E^M(\cdot)$ is the expectation operator under the probability measure induced by the action process under the managerial beliefs. The agent's objective function includes the present discounted value of compensation and the potential private benefits from taking action $a_t < 1$ (the second term). Under the incentive compatible contract Π, I define the agent's expected payoff at time $t \geq 0$ under her beliefs (perceived continuation value) W_t^M as

$$W_t^M = W_t^M(\Pi) = E_t{}^a \left[\int_t^\tau e^{-\gamma t} dC_t \right] \tag{7}$$

I focus the analysis in the case where it is optimal for the principal to implement $e_t = 1$ all the time and provide a sufficient condition for the optimality of implementing this action in Appendix A. Henceforth, the expectation operator $E(\cdot)$ is under the measure induced by $\{e_t = 1 : 0 \leq t < \tau\}$, unless otherwise stated. I define a contract Π incentive compatible if it implements the efficient action.

2.2.2 Principal Problem

At the time the contract is initiated, the firm initial values A_0 and K_0. Given an initial payoff of W_0 for the agent, the principal receives the firm cash flows net of investment cost and pays the compensation to the manager. As a result, given the contract Π, the principal's expected payoff can be written as

$$\hat{P}(W,A,K) \equiv E\left[\int_0^\tau e^{-rt}\left(dX_t - A_t K_t \mathcal{C}(i_t)dt - dC_t\right) + e^{-r\tau} R A_\tau K_\tau \;\middle|\; W_0 = W, A_0 = A, K_0 = K\right]. \tag{8}$$

The objective of the principal is to maximize the present value of the firm cash flows plus termination value net of the agent's compensation, where we make the usual assumption that the principal possesses full bargaining power. The investor's optimization problem reads

$$P(W,A,K) \equiv \max_{\Pi \in IC} \hat{P}(W,A,K) \quad \text{s.t. } W_t \geq 0 \text{ and } dC_t \geq 0 \quad \forall t \geq 0. \tag{9}$$

I denote by $\Pi \equiv (i, C, \tau)$ the solution to this optimization problem.

3 Model Solution

In this section I solve first the first-best model which boils down to the standard neoclassical model of investment, and then the model with agency conflicts.

3.1 A Neoclassical Benchmark

I begin by characterizing optimal investment in the standard neoclassical setting without an agency problem. This benchmark corresponds to the case in which $\sigma_a = 0$ or $\phi = 0$. Given the stationarity of the economic environment and the homogeneity of the production technology, there exists an optimal investment-capital ratio that maximizes the present value of the firm's cash flows. The scaled first-best value of the firm per unit of $A_t K_t$ is therefore given by

$$q^{FB} = \max_{i \geq 0} \frac{\mu - \mathcal{C}(i)}{r + \delta - \mu - i}, \tag{10}$$

where $\mathcal{C}(i) = \frac{1}{2}\lambda_k i^2$ is the quadratic investment cost and the usual conditions for the problem to be well defined hold. I also assume throughout that the firm is sufficiently productive that liquidation is not efficient, that is, $q^{FB} > R$.

From the first-order condition for (10), first-best investment satisfies the usual relation in q-theory

$$\mathcal{C}'(i^{FB}) = q^{FB} = \frac{\mu - \mathcal{C}(i^{FB})}{r + \delta - \mu - i^{FB}}. \tag{11}$$

Equation (11) is the investment Euler equation: the firm invests until the marginal cost of investment equals the marginal value of an additional unit. Because adjustment costs are convex, first-best investment is increasing in q^{FB}. Provided that an interior solution exists, the economically relevant solution for the problem is

$$i^{FB} = (r + \delta - \mu) - \sqrt{(r + \delta - \mu)^2 - \frac{2\mu}{\lambda_k}}, \qquad q^{FB} = \lambda_k \, i^{FB}. \tag{12}$$

Since the manager is more impatient than the principal ($\gamma > r$), it is optimal under symmetric information to pay the manager immediately in cash. As a result, the principal's value function takes the form

$$p^{FB}(w) = q^{FB} - w, \tag{13}$$

where $w = W/(A_t K_t)$ is the manager's scaled continuation utility. In the neoclassical setting, the first-best investment rate is time invariant, and independent of firm's history or volatility of its cash flows.

3.2 Model with Agency Conflicts

I now solve for the optimal contract when there is an agency problem. Recall that the contract specifies the firm's investment policy I, payments to the agent C, and a termination date τ all as functions of the firm's profit history. The contract must be incentive compatible (that is, induce the agent to choose $e_t = 1$ for all t) and maximize investors' value function $P(W, A, K)$[6]. Given an incentive compatible contract and the history up to time t, the discounted expected value of the agent's future compensation is given by ψ. As DeMarzo and Sannikov, 2006 or DeMarzo et al., 2012, I use the martingale representation theorem to show that the perceived continuation payoff of the agent solves

$$dW_t^M = \gamma W_t^M dt - dC_t + \frac{\psi_t}{\sigma_a(1+\beta)} K_t(dA_t - (1+\alpha)\mu A_t dt) \tag{14}$$

This equation shows that the agent's perceived continuation value must grow at a rate to compensate for her time preference. In addition, compensation must be sufficiently sensitive to firm performance, as captured by the processes, to maintain incentive compatibility. Equating the marginal cost of stealing (low effort $e_t = 0$) with the marginal benefit of not stealing (full effort $e_t = 1$) I obtain the Incentive Compatible Constraint

$$\psi_t \geq \frac{\phi\sigma_a(1+\beta)}{\mu(1+\alpha)} = \lambda^{Biased}\sigma_a \tag{15}$$

where $\lambda^{bias} = \frac{\phi(1+\beta)}{\mu(1+\alpha)}$. Note that an optimistic ($\alpha > 0$) and overconfident ($\beta < 0$) will be cheaper to incentivize.

Under the first-order approach, the state variable of the agent is taken into account by the principal to find a necessary condition to satisfy the manager incentive compatibility contraint (ICC) to implement the full effort $e_t = 1$. Afterwards, the state variable that

[6] I restrict attention to regular performance-based contracts and exclude side bets that exploit disagreement about null events. This rules out arbitrage-like behavior.

determines the investor problem is based on the true process A_t and evolves according to

$$dW_t = \gamma W_t dt - dC_t + \frac{\psi_t}{\sigma_a} K_t (dA_t - \mu A_t dt) \tag{16}$$

where ψ_t satisfied the ICC before mentioned. Once incentive compatibility is ensured, the principal's continuation problem can be written under her measure, with managerial beliefs entering only through incentive compatibility (as in De Marzo and Sannikov, 2017 and Rivera, 2023). [7]

The investor's value function in an optimal contract, given by $P(W_t, A_t, K_t)$, is the highest expected payoff the investor may obtain given A_t, K_t, and W_t. While there are three state variables in our model, the scale invariance of the firm's environment allows us to write $P(W_t, A_t, K_t) = A_t K_t p(w_t)$ and reduce the problem to a single state variable: $w_t \equiv \frac{W_t}{A_t K_t}$, the scaled promised payments to the agent as in He, 2009 and DeMarzo et al., 2012.

To characterize the optimal compensation policy and its effects on the investor's (scaled) value function p(w), note that it is always possible to compensate the agent with cash so that it costs at most 1 dollar to increase w by dollar 1 and $p''(w) \geq 1$. In addition, as shown by Eq. (7), deferring compensation increases the growth rate of W (and of w) and thus lowers the risk of liquidation but is costly due to the agent's impatience, $\gamma \geq r$. As a result, the optimal contract sets $dc \equiv \frac{dC}{K}$ to zero for low values of w and only stipulates payments to the manager once the firm has accumulated sufficient slack. That is, there exists a threshold $\overline{w}$ with: $p'(\overline{w}) = -1$ and $dC = max\{0, w - \overline{w}\}$ where the optimal payout boundary is determined by the super-contact condition: $p''(\overline{w}) = 0$. When w falls to zero, the contract is terminated and the firm is liquidated so that: $p(0) = R$. When $w[0; \overline{w}]$, the agent's compensation is deferred and $dc = 0$. The Hamilton–Jacobi–Bellman equation for

[7] Since the principal knows the true process and the parameters governing managerial beliefs (α and β) are static, they only affect the ICC constraint in the principal problem. Hence the HJB does not require an additional state variable tracking the disagreement. If the beliefs of one of the economic agents were to change over time, the state space would generally have to be enlarged. These exercises (as in Miao and Rivera, 2016, Epstein and Ji, 2013...) are outside of the scope of this paper.

the principal's problem is then given by:

$$rp(w) = [\mu - c(i)] + p(w)[\mu + i - \delta_k] + wp'(w)[\gamma - (i - \delta_k + \mu)] + \frac{1}{2}p''(w)\sigma_a^2(w - \lambda^{Biased})^2 \quad (17)$$

subject to the incentive compatibility constraint. Due to the scale invariance (i.e., $P(W, A_0, K_0) = p(w)A_0K_0$) the investor's maximization problem at $t = 0$ can now be rewritten as: $\max_{w_0 \in [0,\overline{w}]} p(w_0)A_0K_0$ with unique solution $w_0 = w^*$ satisfying $p'(w^*) = 0$.

The main results about the optimal contract are that the scaled firm value $p(w)$ is the unique solution to the HJB above, subject to the proposed constraint, over $[0, \bar{w}]$. For $w > \bar{w}$, the scaled value function satisfies $p(w) = p(\bar{w} + (w - \bar{w}))$. Scaled cash payments $dc = \frac{dC}{AK}$ reflect w back to $\bar{w}$. Moreover, the function $p(w)$ is strictly concave on $[0, w]$. I illustrate the proof in Appendix A.

Using the HJB and doing a FOC, I obtain the optimal investment rate:

$$i(w) = \frac{p(w) - p'(w)w}{\lambda_k} \quad (18)$$

It is worth noticing that managerial beliefs will affect the optimal investment rate by acting on the dynamics of the key state variable (scaled deferred compensation). Intuitively, managerial biases will make more or less expensive incentive compatibility and will affect the firms' available resources for investment.

I apply the ITO's lemma and I report the dynamics of the key state variable of the HJB:

$$dw_t = \left[w_t(\gamma - \mu - i + \delta_k + \sigma_a^2) - \lambda^{biased}\sigma_a^2\right] dt + \sigma_a(\lambda^{biased} - w_t)dZ_t \quad (19)$$

It is worth highlighting that, as in He, 2009, there are absorbing states. These first-best absorbing states are attained when the agent has a long history of successes, or equivalently, when the firm has experienced rapid growth. Moreover, behavioural biases α and β will affect the level absorbing state.

3.3 Economic Discussion

Figure 1 reports comparative statics of the model for different values of the belief parameter α[8]. For ease of exposition, I display the two main objects: the principal's firm value $p(w)$ and the investment rate $i(w)$. Recall that the incentive compatibility constraint implies that the higher α, the less binding the constraint becomes, as an optimistic manager is cheaper to incentivize.

[Insert Figure 1 about here]

The first panel, which plots the firm value function $p(w)$ across levels of continuation utility w, shows that higher values of α (optimism) shift the entire value function upward. This reflects the fact that optimism relaxes the incentive compatibility constraint: for a given promised utility, the principal can extract more surplus, as the optimistic manager overestimates the expected payoff of future states. Consequently, the firm's value increases from the principal's perspective. Importantly, since the principal is less concerned about termination, she compensates the optimistic manager earlier, leading to a faster payout schedule and a higher weight on short-term compensation. Hence, beliefs affect not only the level but also the composition and timing of executive pay.

The second panel illustrates the effect on corporate investment. When the manager is optimistic, the first-best outcome is attained for lower values of w, implying that less deferred compensation is needed to implement efficient effort. In this region, the firm invests more aggressively as the cost of providing incentives declines. Relaxing the incentive constraint raises the equilibrium investment rate $\ell(w)$. The shape of the investment curve confirms this intuition: for optimistic managers ($\alpha > 0$), the investment function lies systematically above

[8] I report the model comparative statics only for changes in the parameter governing the bias on the first moment (α). The incentive compatibility condition shows that the parameters governing managerial beliefs α and β rescale the ICC proportionally to $\frac{(1+\alpha)}{(1+\beta)}$. Therefore, it is immediate to see how a bias on the second moment affects the theoretical results in this section.

the benchmark, while for pessimistic ones ($\alpha < 0$) it lies below. This result highlights how small differences in beliefs can potentially affect investment decisions.

Figure 2 quantifies the welfare implications of belief heterogeneity by comparing the percentage change in firm value relative to the rational benchmark. The left panel shows the case of a pessimistic manager ($\alpha < 0$), while the right panel reports the optimistic case ($\alpha > 0$).

[Insert Figure 2 about here]

The left panel highlights that, under pessimism, the cost of providing incentives is significantly higher. When w is low, the firm's total value is small and the cost of motivating effort is limited. As w rises, the constraint tightens and the marginal cost of incentives peaks, before eventually declining again in high states as the manager's continuation utility dominates the termination concern. The resulting curve exhibits a U-shape, with the largest efficiency loss around intermediate states where the agency friction binds most strongly.

The right panel, corresponding to an optimistic manager, displays the mirror image. In this case, the incentive constraint is weaker and the cost of providing incentives is lower throughout the state space. Firm value rises relative to the rational benchmark, with gains peaking at intermediate levels of w, where optimism provides the greatest relaxation of the constraint, and tapering off in both extremes. Quantitatively, the maximum gain in firm value is around 1.4%, while the corresponding loss under pessimism reaches about 1.7%. The symmetry of the two panels underscores how belief distortions act as a wedge that redistributes firm resources between the principal and the agent: optimism benefits shareholders through lower incentive costs, whereas pessimism increases agency rents and reduces firm value. A similar argument could be made for managers that display overconfidence and underconfidence.

Overall, these comparative statics clarify the central mechanism of the model: managerial beliefs affect the incentive compatibility constraint, thereby affecting compensation design, investment behaviour, and firm value. The resulting equilibrium captures a rich interaction between behavioural biases and dynamic contracting that can be quantitatively evaluated using the estimation framework developed in the next sections.

4 Data and Measurements

I use firm accounting and compensation variables from WRDS Compustat Annual and Execucomp dataset from 2005 to 2019. All the data are deflated and expressed in US dollars 2000, with the GDP deflator by the St.Louis Fed. All the variables are winsorized at the 1th and 99th percentile. I exclude financial services firms (SIC codes 6000 to 6999), Utilities (SIC codes 4900 to 4999), Regulated (SIC 8000 to 9999). I drop firms with missing values or values below 5mln of total assets or sales, or with PPE less than 5mln, as standard in the literature.

I define total capital as in Peters and L. Taylor, 2017. Total capital is the sum of Physical capital and Intangible capital. The latter is defined as the sum of Organizational capital and Knowledge capital. The stock of Organizational capital and Knowledge capital is computed with the perpetual inventory method on SGA and R&D with a depreciation rate of 30% and 15%, respectively. Similarly, intangible investment are defined as R&D plus 30% SGA. We can see below the description for variable construction and the summary statistics. I use the dataset execucomp for managerial compensation. The main variable of interest will be TDC1 (salary, bonus, other annual, restricted grant, stock options, long term incentives and all other compensation...). In particular, the measure of deferred compensation considered is TDC1 - (salary + annual bonus + other annual).

A key part of this work revolves around quantifying the effect of managerial beliefs. I use managerial guidance about future sales to proxy for managerial expectations (as in C. Otto, 2014). There are mainly two considerations when using these data (Ma et al., 2024): manipulation and panel size, particularly for managerial guidance. I rely on Factset and IBES to extend the coverage and collect this data. Moreover, I restrict the disclosure to be only in the short-term (12 months before the fiscal year). The reasoning is that if managers manipulate sales guidance, they will do it more in the long-term as in the short-term, they cannot go too far from what they truly believe without avoiding consequences. The reported expectations of future sales come with a range of sales, and this allows to study both the mean and variance of the subjective distribution. Moreover, managers report forecasts for various items and after a careful analysis, I conclude that sales is the best one for this study. I provide a detailed explanation of the data cleaning and collection process in Appendix B. The final sample includes 1,704 firms-year observations about firms from 2005 to 2019.

[Insert Table 1 about here]

The table above describes the summary statistics of the final sample. The sample is slightly tilted towards firms with high valuation and larger ones.

5 Stylized Facts

This section documents a set of empirical regularities on managerial expectations, forecast accuracy, and compensation that motivate and inform the structural analysis developed in later sections. The goal is to establish that managerial forecasts are both informative and systematically biased in ways consistent with the belief distortions modelled in the theory. I begin by assessing whether managerial sales forecasts contain predictive information about realised firm outcomes. I then characterise the cross-sectional and time-series distribution of forecast errors and their variances, highlighting heterogeneity in optimism and confidence across managers and over time. Finally, I examine how these belief distortions correlate with the structure of executive compensation.

5.1 Forecast Accuracy

Before using managerial forecasts to estimate belief distortions, it is essential to verify that they contain meaningful information about future firm performance. This section assesses the informational content of managerial expectations by examining whether forecasted sales predict realized outcomes. Following the standard approach in the expectations literature, I estimate the following regression:

$$Sales_{i,t} = \alpha + \beta F_{i,t-1}\big[Sales_{i,t}\big] + \epsilon_{i,t}, \tag{20}$$

where $F_{i,t-1}$ denotes the forecast of next-period sales made by manager i at time $t-1$. Both realized and forecasted variables are scaled by lagged total capital to ensure comparability across firms and to mitigate the influence of firm size. All specifications include year fixed effects, and column (2) additionally absorbs industry fixed effects. Standard errors are clustered by year in column (1) and two-way clustered by industry and year in column (2).

[Insert Table 2 about here]

Table 2 shows that managerial forecasts of sales relative to assets are highly accurate predictors of realized outcomes. The estimated slope coefficient is 0.91 in the specification with year fixed effects only, and remains 0.91 once industry fixed effects are added (both significant at the 1% level), indicating a very tight linear relationship between expected and realized sales. The adjusted R^2 in both specifications confirms that most of the variation in realized sales is explained by managerial expectations. This strong predictive power indicates that forecasts are highly informative and closely aligned with actual firm fundamentals, rather than being mere noise or mechanical extrapolations. Taken together, these results support the interpretation of managerial forecasts as credible measures of subjective expectations about firm performance. They validate their use in the structural estimation of belief distortions in the next sections.

5.2 Compensation

The goal of this section is to test model implication about the relationship between managerial beliefs and their incentive compensation, I estimate the following empirical specification:

$$Incentives_{i,t} = \alpha + \beta_1 ExcessSubjectiveMean_{i,t} + \beta_2 ExcessSubjectiveVol_{i,t} + \epsilon_{i,t}, \quad (21)$$

where $Incentives_{i,t}$ measures the value of deferred compensation pay. The development of these measures follow Barrero, 2022. In particular, $ExcessSubjectiveMean_{i,t}$ denotes the manager's subjective excess mean of sales and is compute by the difference between the mid point of sales forecasts and the realization. A positive value indicates optimism and a negative value suggests pessimism. $ExcessSubjectiveVol_{i,t}$ captures the manager excessive subjective variance of sales. In particular, the excess subjective volatility is computed as the opposite of the Excess Absolute Forecast Error[9]. All the values are scaled by lagged total capital (Peters and L. Taylor, 2017). A positive value means that managers are overstating uncertainty (underconfident) and a negative one means that the subjective uncertainty is lower than the realized one (overconfident).

The model predicts that β_1 should be negative: optimistic managers (who systematically overestimate future sales and thus display positive forecast errors) are easier to motivate and require lower-powered incentives. Conversely, β_2 is expected to be positive, as underconfident managers (who overestimate uncertainty and exhibit higher excess forecast variance) are harder to incentivize and therefore receive stronger pay-for-performance schemes. This regression provides a direct empirical link between observed belief distortions and the structure of executive compensation implied by the theory.

[Insert Table 3 about here]

Table 3 provides empirical support for the theoretical predictions. The coefficient on forecast error ($ExcessSubjectiveMean$) is consistently negative and highly significant across specifications, indicating that optimistic managers—those who overestimate future firm performance—receive weaker incentive pay relative to their less optimistic peers. This finding is consistent with the model's prediction that optimism relaxes the incentive compatibility constraint, allowing the principal to economize on costly deferred compensation.

By contrast, the coefficient on the excess variance of forecast errors ($ExcessSubjectiveVol$) is positive and significant, confirming that underconfident managers—those who perceive higher uncertainty—require stronger incentive intensity to align their effort with shareholder objectives.

Overall, these results reinforce the model's central mechanism: optimism reduces, while

[9]The excess absolute forecast error is defined closely following Barrero, 2022, it is the difference between the firm's actual absolute forecast error for sales and its subjective absolute forecast error or subjective uncertainty)

underconfidence increases, the cost of providing managerial incentives. They suggest that observed differences in executive pay structures can be rationalized as equilibrium responses to heterogeneous managerial beliefs.

6 Estimation

The structural estimation aims to recover the deep parameters governing firm dynamics and managerial beliefs from observable data. The key challenge is that the key state variable in the model (scaled continuation utility) is not directly observable, while the available firm-level data are noisy, incomplete, and often measured with error. To address these issues, I cast the model into a state-space representation that links unobserved latent states to observable firm-level quantities such as sales, investment, and forecasts. This formulation enables consistent inference under measurement error and missing data. The parameters are then estimated by maximum likelihood using an Unscented Kalman Filter (UKF), which efficiently handles the model's nonlinearity.

6.1 State Space and Measurement Equations

Plugging the optimal firm's policies in the dynamics of the state variables gives the controlled dynamics of these processes. The result is the optimal time series trajectory of each controlled variable, which is therefore free of any endogeneity bias. In fact, the model describes how each endogenous variable determines observable (possibly noisy) quantities like capital through exogenous model parameters, which are the objective of the inference procedure. The key equations of interest are:

$$\begin{cases} K_{t+1} = K_t + (i_t - \delta_k)K_t \\ W_{t+1} = (1+\gamma)W_t + \lambda_{biased}\sigma_x A_t K_t \epsilon^x_{t+1} \\ A_{t+1} = A_t(1 + \mu + \sigma_a \epsilon_{t+1}) \end{cases} \tag{22}$$

I will consider $Z_t = A_t K_t$ as a state variable to simplify the implementation. After applying Ito's lemma, I get:

$$dZ_t = (\mu + i_t - \delta_k)Z_t dt + \sigma_a Z_t dW_t \tag{23}$$

which describes its dynamics. Therefore the key state variables will be $X_t = [W_t; Z_t]$. I can write $X_{t+1} = f(X_t; u_{t+1})$ which is therefore M(1) where $u_{t+1} \sim \mathcal{N}(0, Q_t)$[10].

Estimation of equations above faces several challenges. First, any period's sales are simultaneously hit by a shock that will have to be identified. Second, the data available to the econometrician are subject to measurement errors. Unaddressed, this error-in-variables problem would result in inconsistent estimates of the model's parameters. Third, Compustat data are plagued by missing values.

Given these setup, the most efficient estimation procedure while addressing these problems is by maximum likelihood with an unscented Kalman filter. In what follows, I describe the steps of this procedure for the case of a complete data set. Appendix B provides the full details, including the case in which there are missing observations. The measurement

[10] $Q_{t+1} = \begin{bmatrix} (\lambda_{biased}\sigma_x Z_t)^2 & \lambda_{biased}(\sigma_x Z_t)^2 \\ \lambda_{biased}(\sigma_x Z_t)^2 & (\sigma_x Z_t)^2 \end{bmatrix}$

equations are given by:

$$\begin{cases} Y_t^* = Y_t + \epsilon_{t,2} = \mu Z_t + v_{1,j,t} \\ \ell_t^* = \ell\left(\frac{W_t}{Z_t}\right) + \epsilon_{t,2} \\ W_t^* = W_t + \epsilon_{t,3} \\ \tilde{E}[Y_t]^* = (1+\alpha)\mu Z_t + \epsilon_{t,4} \\ [\tilde{V}[Y_t] - V[Y_t]]^* = \beta(2+\beta)(\sigma_x Z_t)^2 + \epsilon_{t,5} \end{cases} \tag{24}$$

where $z_{1,j,t}, \ldots, z_{5,j,t}$ are, respectively, the noisily observed operating earnings, capital stock, labor stock, investment and hiring of firm j in year t. The measurement errors, $v_{1,j,t}, \ldots, v_{5,j,t}$ have variances $\sigma^2_{v,1}, \ldots, \sigma^2_{v,5}$. I let this set of equations and parameters represent a set of firms $j = 1, \ldots, N$. Therefore, the vector of measurement errors for each group of N firms, $\mathbf{v}_t$, is $5N$-dimensional with $\mathbf{v}_t \sim \mathcal{N}(0, \mathbf{R})$ and $\mathbf{R}$ a diagonal covariance matrix.

Altogether, this state space model will be used to identify four parameters: λ_k, σ_a, α, β and the variances of the measurement errors $\sigma^2_{v,1}, \ldots, \sigma^2_{v,5}$; while the other parameters will be calibrated (r=0.1, $\gamma = 0.12$, $\delta_k = 0.2$, $R = 0.33$, $\mu = 0.07$, $\phi = 0.4$). The measurement equation capturing the forecast error will be identifying the parameter α; the difference in the squares of forecast and realization will help identifying the parameter β [11] Estimation at the firm level is infeasible because the time series are too short. To achieve high granularity I follow the same approach as in Gryglewicz et al., 2021, Frésard et al., 2023, Gashi, 2025 to estimate the model's parameters for each of many small groups of very similar firms, namely $N = 20$, assuming each firm in the group has the same parameter. I will explain this in details in the next section.

It is also worth discussing that in the model W_t is a stock, while in the data I observe only flows. As discussed in Appendix C, I convert the flow to the stock with minimal

[11] As we can see from the measurement equation, this is a second-order equation and therefore two solutions satisfy it. Economic reasonig makes only one β acceptable.

assumption as Ward, 2025, He, 2011, Ai et al., 2023, Gryglewicz and Hartman-Glaser, 2020. Given the model parameters, the unscented Kalman filter recovers the unobserved state process X_t' that determines the likelihood function of the observed $5N$-dimensional data $z_t \equiv [z_{1,1,t}, \ldots, z_{5,j,t}, \ldots, z_{1,40,t}, \ldots, z_{5,40,t}]$, i.e., earnings, capital, labor, investment and hiring of the N firms in each group and for $t = 1, \ldots, T$:

$$\sum_{t=1}^{T} -\frac{1}{2}\Big[5N\log(2\pi) + \log|F_{t|t-1}| + (z_t - \hat{z}_{t|t-1})' F_{t|t-1}^{-1}(z_t - \hat{z}_{t|t-1})\Big], \tag{25}$$

where $\hat{z}_{t|t-1}$ is the one-step-ahead prediction of z_t based on the filtered state process x_t, and $F_{t|t-1}$ is the error covariance matrix.

6.2 Granularity and Grouping Mechanisms

Ideally, the structural model would be estimated at the individual firm level. However, given the computational complexity of the estimation procedure and the limited time-series dimension for each firm, I instead estimate the model described in Section 2 on groups of firms observed over a fixed time window. This approach balances statistical precision and computational feasibility while preserving meaningful cross-sectional variation in parameters.

To determine the appropriate level of aggregation, I perform a series of diagnostic analyses to assess small-sample performance and estimation stability. Groups of approximately twenty firms provide a suitable compromise, ensuring sufficient within-group observations for reliable estimation while retaining heterogeneity across groups. Within each group, firms are assumed to share the same deep structural parameters, such that differences across groups capture heterogeneity in beliefs and technologies.

Grouping follows three economically motivated dimensions: (i) industry classification, (ii) managerial subjective mean, and (iii) excess subjective variance, following the methodology of Frésard et al., 2023 and Gryglewicz et al., 2021. First, I partition firms by Fama–French five-industry classification to account for common exposure to sector-specific shocks such as labor shortages or technological disruptions. Second, I sort firms by the average forecast error of their managers, capturing heterogeneity in optimism or pessimism (the α parameter). Third, I sort by the excess variance of forecast errors, which proxies for differences in managerial confidence (the β parameter). Applying this grouping procedure yields 18 distinct groups covering a total of 405 firms, providing a rich cross-section for identifying belief and incentive heterogeneity.

7 Results

This section presents the paper's main quantitative findings. I begin by reporting structural parameter estimates obtained from group-level estimation (Section 5) and discuss their economic interpretation. I then show how the estimated belief parameters map into industry-level patterns in expectations and incentive intensity, and I relate these patterns to observable compensation schemes. Finally, I quantify the role of belief heterogeneity through a counterfactual experiment that compares observed outcomes to a rational-expectations benchmark. Together, these results illustrate how subjective beliefs shape contracting and corporate decisions.

7.1 Estimates

Table 4 presents the main parameter estimates from the structural model. For each parameter, I report the mean of the group-level point estimates, their cross-group standard deviation, and the percentiles of the distribution (p25, p50, p75). The parameters are estimated using the procedure described in Section 5, combining the Unscented Kalman Filter with simulated maximum likelihood on industry-homogeneous groups of firms. Overall, the estimates are economically plausible and consistent with the mechanisms discussed in Section 3.

[Insert Table 4 about here]

The average estimate of λ_k around 1.67 suggests that the marginal cost of investing is more than linear and its dispersion shows meaningful heterogeneity across industries.

The estimated volatility parameter σ (mean 0.084) indicates moderate fluctuations in productivity. This parameter governs the uncertainty related to cash flows in the model and therefore the intensity of the moral hazard problem in this setup.

Turning to the parameters of managerial beliefs, α and β capture the degree of optimism (first-moment bias) and confidence (second-moment bias), respectively. The mean estimate of α is -0.009, indicating a mild pessimistic bias on average, though the interquartile range (-0.015 to 0.009) still straddles zero, so optimistic and pessimistic managers coexist within the sample. The picture is markedly different for β: with a mean of -0.006 and an interquartile range running from -0.006 to -0.002, the estimates are negative across essentially the entire distribution. This points to a systematic overprecision bias that is, managers tend to underestimate the variance of firm-level shocks.

Table 5 reports the distribution of the absolute value of the t-stats for the estimated parameters. On average, the precision of the estimates is high. This reflects the effectiveness of the Unscented Kalman Filter in extracting information from noisy and incomplete data.

[Insert Table 5 about here]

Overall, the estimates confirm that managerial belief distortions are statistically significant and economically meaningful despite being small in magnitude. The results provide quantitative support for the model's central mechanism: differences in optimism and confidence alter the perceived continuation value of the contract, thereby shaping the cost of incentives and the allocation of resources across firms. This cross-sectional variation will be further exploited in the counterfactual analysis that follows, where I examine how firm value and investment would change under rational expectations.

7.2 Model Fit

The estimation results reveal substantial heterogeneity in belief distortions across industries. However, a key question remains: are these belief parameters statistically necessary to explain the data, or would a simpler model with fully rational managers provide an equally good fit? This question is crucial to confirm that the additional model complexity is needed to provide an accurate representation of reality.

To address this question, I employ likelihood ratio (LR) tests to formally compare the unrestricted model against three nested specifications. The first restricts the optimism parameter to zero ($\alpha = 0$), implying that managers have unbiased expectations about the first moment and there is no optimism/pessimism. The second restricts the parameter β about the second moment to zero, implying that managers correctly assess the variance of cash flows and exhibits no overconfidence or underconfidence. The third imposes both restrictions simultaneously ($\alpha = \beta = 0$), corresponding to the standard rational expectations benchmark commonly assumed in the agency literature.

For each group, I compute the LR statistic as twice the difference between the unrestricted and restricted log-likelihoods. Under standard conditions, this statistic follows a chi-squared distribution with degrees of freedom dependingg on the number of restrictions. I then compare each statistic against the appropriate critical value at the 5% significance level and report the fraction of groups in which the null hypothesis is rejected.[12]

[Insert Table 6 about here]

Table 6 presents the results. The hypothesis of no optimism bias is rejected in approximately 58% of industries, while the hypothesis of correct confidence is rejected in about 63% of industries. Most importantly, the joint hypothesis of managerial rationality is rejected in nearly three-quarters of the sample. These rejection rates indicate that the belief distortions documented in this paper are not merely point estimates that happen to differ from zero, they represent statistically significant departures from the rational benchmark in the majority of industries. Taken together, these results validate the core premise of this paper: incorporating belief heterogeneity meaningfully improves our ability to explain observed patterns in managerial forecasts and firm investment, and the standard rational expectations assumption is rejected by the data.

[12] I also conduct pooled tests that aggregate the LR statistics across all groups. All three null hypotheses are rejected with p-values below 0.001, providing additional support for the unrestricted model at the aggregate level.

7.3 Industry

I next examine how managerial beliefs are heterogeneous and their impact on incentive structures across industries.

7.3.1 Optimism and Pessimism

Figure 3 compares the estimated degree of optimism and pessimism across industries. In the model, belief distortions in the first moment are captured by the parameter α, which shifts the manager's subjective expectations of future cash flows. A positive α indicates optimism: the manager expects higher future sales than the principal, perceiving projects as more profitable. Conversely, a negative α reflects pessimism: the manager expects lower future outcomes than the principal, perceiving projects as less attractive. The parameter α therefore provides a direct measure of how the perceived profitability deviates from the rational benchmark.

Belief distortions of this kind have first-order implications for contracting. Optimism relaxes the incentive compatibility constraint: an optimistic manager attaches greater value to continuation utility and is therefore willing to exert effort under weaker pay-performance sensitivity and lower deferred compensation. Pessimism operates in the opposite direction, tightening the constraint and forcing the principal toward steeper incentive schemes to induce comparable effort.

Figure 3 documents pronounced heterogeneity in α across industries. Consumer Goods, Manufacturing, and the residual "Other" category display positive point estimates, with means of approximately 0.009, 0.008, and 0.018 respectively, indicating that managers in these sectors hold mildly optimistic beliefs relative to the rational benchmark. Healthcare ($\alpha = -0.002$) is essentially indistinguishable from zero. The most striking departure is observed in Technology, where the estimated mean $\alpha of -0.032$ is negative, pointing to systematic managerial pessimism.

[Insert Figure 3 about here]

Taken together, these estimates indicate that the direction and magnitude of first-moment belief distortions vary systematically with the economic environment in which firms operate.

7.3.2 Overconfidence and Underconfidence

Figure 4 illustrates the role of the parameter β, which captures the manager's perception of uncertainty. A negative β corresponds to overconfidence: the manager underestimates the variance of outcomes. Conversely, a positive β corresponds to underconfidence: the manager overestimates uncertainty.

These biases influence firm value and investment via the optimal contract. Overconfident managers will be cheaper to incentives as they will understate uncertainty and less incentives compensation will induce full effort provision. The opposite is true for the underconfident manager, who will demand higher pay-performance sensitivity. The empirical estimates reveal a clear and systematic pattern across industries. Figure 4 shows that the average β is *negative* in every industry considered, indicating that managerial overconfidence, the tendency to underestimate the volatility of firm-level shocks relative to the rational benchmark, is a pervasive feature of the cross section rather than a sector-specific phenomenon. The magnitude of this bias, however, varies substantially across industries. Manufacturing exhibits the most pronounced overconfidence ($\beta = -0.013$) followed by Healthcare ($\beta = -0.006$) and Technology ($\beta = -0.005$). Consumer Goods ($\beta = -0.003$) displays rational like behaviour, and the "Other" category lies closest to the rational benchmark.

These patterns have direct implications for incentive provision. Because overconfident managers underestimate the variance of cash-flow innovations, the principal can therefore implement effort with lower deferred compensation than would be optimal under rational beliefs, reducing the agency costs.

[Insert Figure 4 about here]

Taken together, the industry-level differences in the first- and second-moment biases underscore the importance of jointly accounting for both dimensions of belief distortion to understand how managerial beliefs shape incentive provision under the optimal contract.

7.3.3 Incentive Provision Across Industries

I now turn to studying the implications of the estimated biases for incentives. Figure 5 examines how incentive provision varies across industries as a function of managerial beliefs. The figure plots the belief-adjusted intensity of incentives required to implement the first-best action under different combinations of optimism (α) and confidence (β). As derived in the model, the effective strength of incentives can be expressed as:

$$\lambda \propto \frac{1+\beta}{1+\alpha}. \tag{26}$$

This expression captures how first- and second-moment distortions jointly determine the tightness of the incentive compatibility constraint. When the manager is fully rational ($\alpha = \beta = 0$), $\lambda = 1$, meaning that the contract implements the first-best with the standard level of incentive intensity. Values of $\lambda > 1$ indicate that stronger incentives are required to induce effort provision equivalently, that managerial beliefs make the agent harder to motivate, while $\lambda < 1$ implies that fewer incentives are sufficient because the manager's beliefs already align more closely with the principal's objectives.

The empirical estimates reveal substantial cross-industry variation in λ, the average value lies *above* the rational benchmark. Consumer Goods ($\lambda = 0.988$), Manufacturing ($\lambda = 0.980$), and the residual "Other" category ($\lambda = 0.982$) all exhibit λ slightly below one, indicating that managerial beliefs in these sectors collectively reduce the cost of incentive provision relative to the rational benchmark. Healthcare lies essentially at the benchmark ($\lambda = 0.996$), with a tight confidence interval. The clear outlier is Technology, where $\lambda = 1.030$ and significantly exceeds one: in this sector alone, belief distortions raise the cost of contracting, requiring shareholders to provide higher-powered incentives to elicit effort.

The sources of these deviations, however, differ across sectors. In Consumer Goods, Manufacturing, and the residual "Other" category, mild first-moment optimism ($\alpha > 0$) and moderate second-moment overconfidence ($\beta < 0$) reinforce each other in the same direction: optimism raises the manager's perceived continuation value, while overconfidence reduces the perceived risk of incentive contracts, jointly relaxing the incentive compatibility constraint and lowering the equilibrium cost of motivating effort. The Manufacturing case is particularly instructive: although it exhibits the strongest overconfidence in the cross section ($\beta = -0.013$), the relatively modest optimism ($\alpha = 0.008$) anchors λ close to the benchmark, because the two biases enter the incentive intensity multiplicatively and partially offset one another. Healthcare displays a near-symmetric configuration in which mild pessimism ($\alpha = -0.002$) is approximately offset by mild overconfidence ($\beta = -0.006$), yielding an incentive intensity similar from the rational benchmark. Technology stands apart precisely because the two biases now compound rather than offset: managers are systematically pessimistic ($\alpha = -0.032$) and only moderately overconfident ($\beta = -0.005$), so the negative first-moment bias dominates the negative second-moment bias and pushes λ above one. Pessimism lowers the manager's perceived continuation value and tightens the participation and incentive compatibility constraints, while the comparatively modest overconfidence is insufficient to offset this effect.

The overall pattern underscores the importance of considering both dimensions of belief heterogeneity jointly. Analyzing optimism and confidence in isolation would misrepresent the true cost of incentive provision, as the two biases can offset or amplify each other depending on their relative magnitudes. For example, a strongly pessimistic yet only mildly overconfident manager—as in Technology—may be substantially more costly to motivate than a mildly optimistic but highly overconfident one—as in Manufacturing. These results emphasize that optimal contracting under belief heterogeneity is inherently multidimensional, with incentive intensity shaped by how first- and second-moment distortions interact.

[Insert Figure 5 about here]

Taken together, these findings provide a unified picture of how managerial beliefs shape the cost and structure of executive compensation across industries.

7.3.4 The Cross Section of Incentives and Firm Characteristics

The goal of this part is to link the cross section of incentives implied by managerial characteristics to firm characteristics. This would be a first step towards understanding which firms should try to hire managers with different characteristics and which companies are incurring more distortions.

[Insert Table 7 about here]

Table 7 examines how the belief-induced incentive distortion varies with observable firm characteristics. For each characteristic, I partition firms into those below and above the sample mean and compute the average incentive ratio $\lambda^{biased} = (1+\beta)/(1+\alpha)$. This ratio captures how much more (or less) expensive it is to incentivize the biased manager relative to a rational counterpart: values above one indicate that belief distortions increase the cost of incentive provision, while values below one indicate that the principal can exploit the manager's biases to reduce compensation. A notable feature of the new estimates is that several below-mean groups display λ^{biased} below one, indicating that belief distortions are not uniformly costly across firms but instead vary in both magnitude and sign with observable characteristics.

Younger and smaller firms exhibit significantly higher incentive distortions. The difference in λ^{biased} between firms below and above the mean is -0.0064 for age (significant at the 1% level) and -0.0042 for size (significant at the 5% level), indicating that managers at younger and smaller firms are systematically more costly to motivate. This pattern is consistent with the view that early-stage firms attract managers who are more uncertain about firm outcomes, tightening the incentive compatibility constraint. Alternatively, less established firms may lack the governance infrastructure managerial expertise.

Intangibilityis associated with the largest positive distortions in the table, the difference is $+0.0133$. This finding suggests that managers at firms with harder-to-value assets require to put in place more expensive incentives.

Financial constraints are also associated with significantly higher incentive distortions. Firms above the mean of the Hoberg-Maksimovic financing delay index exhibit λ^{biased} of 1.0069, compared to 0.9996 for less constrained firms, a difference of $+0.0073$ (significant at the 1% level). This pattern complements the age and size results: firms that find external finance more difficult to obtain also tend to employ managers whose beliefs depart more sharply from the rational benchmark, pushing up the cost of providing incentives.

A distinct pattern emerges for institutional ownership. Firms with above-mean institutional ownership display λ^{biased} of 1.0003, compared to 1.0119 for firms below the mean, a difference of -0.0116 (significant at the 1% level). This is the largest effect in the table in absolute value and points to a governance and monitoring channel.

Taken together, these cross-sectional patterns paint a coherent picture. Belief-induced incentive distortions are concentrated among younger, smaller, more intangible-intensive, and more financially constrained firms, while stronger institutional monitoring is associated with substantially smaller distortions. These findings complement the industry-level analysis and provide firm-level evidence that belief heterogeneity is not uniformly distributed but instead varies systematically with observable characteristics.

7.4 Counterfactual Analysis

I now use the estimated parameters to quantify the economic importance of the mechanisms highlighted by the theory. The analysis is organized around two questions that speak directly to the normative content of the model. First, how much value do sophisticated boards create by designing contracts that endogenize managerial belief distortions? Second, how costly are managerial biases to shareholders once the optimal contract is in place, and what would be the welfare gains from replacing biased executives with fully rational counterparts? Together, these exercises discipline the quantitative magnitudes implied by the model, map them into economically meaningful corporate outcomes, and clarify why behavioral biases can survive in equilibrium despite being observed at the top of the corporate hierarchy.

7.4.1 Unsophisticated Board and Suboptimal Contracting

Throughout this paper, I have maintained the assumption that the principal designs and implements the optimal contract, fully accounting for the manager's belief distortions. This assumption implies that boards are sophisticated: they accurately know managerial biases and exploit this information to minimize contracting costs while maintaining incentive compatibility. While potentially strong, this assumption finds support in the empirical corporate governance literature, which documents that boards actively design compensation packages and incentives tailored to managerial characteristics.

Nevertheless, this assumption naturally raises a counterfactual question: *What are the consequences when boards fail to account for managerial belief? What is the value of the optimal contract when the principal and the agent disagree?* In other words, what happens when an unsophisticated board offers the manager a contract designed for a rational agent ($\alpha = \beta = 0$) rather than a contract tailored to the manager's actual beliefs? This counterfactual exercise quantifies the value of board sophistication and provides insight into the potential costs of suboptimal contracting in the presence of behavioral biases. The consequences of offering a "rational" contract to a biased manager depend critically on the direction of the belief distortion, captured by the bias ratio $\lambda^{Biased} = (1 + \beta)/(1 + \alpha)$.

Case 1: Cheap Managers ($\lambda^{Biased} \leq 1$). In this case, offering a rational contract (with higher incentive intensity than necessary) still satisfies the incentive compatibility constraint the manager continues to exert effort. The primary consequence is that the principal fails to exploit the manager's optimism: the contract provides stronger incentives than required, transferring surplus to the manager that could otherwise be retained by shareholders. While suboptimal, this scenario does not lead to a breakdown of effort provision. The counterfactual in this case compares firm value and investment under the estimated contract versus under a rational benchmark, quantifying the "exploitation value" of managerial optimism.

Case 2: Expensive Managers ($\lambda^{Biased} > 1$). The agency conflict becomes substantially more severe in this case. Here, the manager requires stronger incentives to exert effort than a rational agent would. If the board, failing to recognize this, offers a contract calibrated for a rational manager, the incentive compatibility constraint no longer binds as the contract provides insufficient incentives from the manager's (biased) perspective. Facing inadequate compensation for effort, the manager optimally chooses to shirk ($e_t = 0$), diverting firm resources for private benefit while providing no productive contribution to cash flow generation. [13]

The results reveal substantial value creation from board sophistication. Across all firms, sophisticated boards that correctly account for managerial biases improve firm value by an average of 13.9 % and investment rates by 11.3 % relative to the unsophisticated benchmark. These aggregate figures, however, mask important heterogeneity across the two cases.

These findings carry important implications for corporate governance. First, they demonstrate that the costs of suboptimal contracting are highly asymmetric. Failing to exploit a cheap manager's optimism results in modest costs; failing to adequately incentivize an expensive manager results in value destruction through shirking and resource diversion. This asymmetry suggests that boards should be particularly vigilant in identifying and accommodating underconfidence or pessimism in managerial beliefs.

Second, the results provide a rationale for the persistence of seemingly generous compensation packages. When managers exhibit underconfidence ($\beta > 0$) or pessimism ($\alpha < 0$), providing "excessive" incentives relative to a rational benchmark is not inefficient but rather necessary to maintain effort provision. Boards that fail to recognize this necessity risk triggering shirking behavior with severe consequences for firm value.

Finally, the results highlight a novel channel through which behavioral biases affect corporate outcomes. Prior literature has emphasized how managerial overconfidence or optimism directly distorts investment decisions. The present analysis reveals an indirect channel operating through the contracting relationship: biases affect the incentive intensity required to maintain effort, and boards that fail to adjust for these biases may inadvertently induce shirking, with consequences that far exceed the direct effects of the biases themselves.

[13] This shirking equilibrium necessitates a distinct theoretical treatment. Under shirking, the cash flow process loses its positive drift: $dA_t = \sigma_a A_t dZ_t$, rather than $dA_t = \mu A_t dt + \sigma_a A_t dZ_t$ under effort. Therefore, the implied behaviour differ than in the model developed in the main section of the paper.

7.4.2 Rational Manager Benchmark

I now quantify the economic cost of managerial belief distortions and assess how much firm value and investment would change if managers were fully rational. This counterfactual exercise helps evaluate whether firms would benefit from replacing biased managers with rational ones.

Formally, I recompute firm values and investment policies under the assumption that all managers hold rational expectations. In practice, I set $\alpha = 0$ and $\beta = 0$ for all groups, thereby aligning the subjective beliefs of managers and shareholders. This isolates the direct impact of managerial beliefs on firm value and corporate investment. The approach parallels the identification logic in Barrero, 2022 and Ma et al., 2024, who similarly contrast observed outcomes with rational-expectations benchmarks.

The main finding is that removing managerial belief distortions has only a modest impact on firms. Investment policies remain virtually unchanged when managers are assumed to be fully rational, indicating that optimal contracts already adjust endogenously to offset differences in beliefs. Firm values do increase under the rational-expectations benchmark, but only by about 0.1% on average. This limited effect contrasts with the larger gains reported in previous studies such as Ma et al., 2024 and Barrero, 2022, reflecting that, in my setting, the optimal contract internalizes disagreement between managers and shareholders. As a result, most of the inefficiencies caused by biased expectations are absorbed within the contracting equilibrium rather than transmitted to real outcomes. This result is quantitatively drivenb by the small sizes of these contracts compared to the overall firms characteristics.

From a normative standpoint, the small welfare gain from removing biases suggests that firing biased managers and replacing them with fully rational ones would yield negligible economic benefits. The potential improvement of 0.1% in firm value represents a gross upper bound, which must be compared to the search, turnover, and adaptation costs associated with managerial replacement. Given that these costs are likely to be substantially larger, it is not surprising that firms tolerate or even accommodate behavioral managers in equilibrium. In this sense, the persistence of optimism and underconfidence can be rationalized in equilibrium with optimal contracting environment.

7.4.3 Benefits and Costs of Biased Managers

The counterfactual analysis presented in Section 7.4.1 suggests that, on average, eliminating managerial biases yields a negligible increase in firm value of approximately 0.1%. However, this aggregate result masks significant cross-sectional heterogeneity. While the optimal contract protects the firm from belief distortions on average, extreme deviations in beliefs can generate meaningful value creation or destruction depending on the direction of the bias.

To explore this heterogeneity, I narrow the analysis to the tails of the distribution. I isolate groups of firms where the impact of managerial beliefs is most pronounced: those where replacing the biased manager with a rational one leads to the largest decrease in firm value (implying the current biased manager creates value for shareholders) and those where it leads to the largest increase (implying the biased manager is costly). I report below the average parameters and counterfactual outcomes for these two extremes. The right column of Table 8 ("Highest $\Delta p(w)$") captures the scenario where belief distortions are most detrimental to shareholders. For these firms, replacing the current manager with a rational agent would increase firm value by 0.06% and raise investment by 0.02%. The structural estimates reveal that these managers are characterized by mild pessimism ($\alpha = -0.02$) together with slight overconfidence ($\beta = -0.01$).

Although the overconfidence component marginally eases the incentive burden by lowering the manager's perceived volatility of the continuation claim, pessimism dominates: the principal must provide stronger, more expensive incentives to compensate for the manager's undervaluation of the firm's prospects. In these cases, the agency conflict is exacerbated by beliefs, resulting in measurable – albeit modest – value destruction.

The left column ("Lowest $\Delta p(w)$") presents the opposite case, where the current biased manager is effectively cheaper than a rational counterpart. Here, switching to a rational manager would actually *decrease* firm value by 0.15% and reduce investment by 0.03%. These managers exhibit mild optimism ($\alpha = 0.02$) combined with slight overconfidence ($\beta = -0.01$). Both biases work in the principal's favor: the optimistic manager overestimates the expected value of their deferred compensation, while overconfidence reduces the perceived volatility of their continuation claim. As a result, the manager is willing to exert effort for a lower promised utility than a rational agent would require. Effectively, the board exploits these biases to extract surplus, paying the manager less in expected terms than the true economic cost of their effort. This result provides a rationale for why boards may prefer biased executives: in specific contexts, managerial optimism and overconfidence act as substitutes for costly monetary incentives.

Overall, Table 8 highlights that the welfare implications of biased beliefs are asymmetric even when the magnitudes are modest. The bias ratio $(1 + \alpha)/(1 + \beta)$ summarizes this intuition: when the ratio exceeds unity (driven by optimism or overconfidence), the manager is "cheap" to incentivize; when it falls below unity (driven by pessimism or underconfidence), the manager becomes expensive, and the agency wedge widens.

[Insert Table 8 about here]

7.4.4 Magnitude and State-Dependency

The baseline counterfactual analysis indicates that, at the current average levels of scaled deferred compensation (the state variable w), the welfare impact of belief distortions is approximately 0.1% of firm value. This modest magnitude reflects the fact that, for the typical firm in the sample, the manager's continuation value constitutes a relatively small fraction of total firm assets[14]. However, a central feature of dynamic contracting is that w is not static; it drifts upward following a sequence of positive shocks. As the manager delivers strong performance, their inside stake grow, potentially affecating the agency friction.

To quantify this state-dependency, I conduct a simulation exercise where I exogenously scale the initial level of the manager's continuation value by a factor $\xi \in [1, 3]$. This effectively simulates the firm's dynamics after a prolonged period of "good news," where the manager's accumulated wealth becomes large relative to the capital stock.

Figure 6 plots the counterfactual gains from removing biases as a function of ξ. The results reveal a strong non-linearity. While the value added from rationality (blue line) rises modestly to approximately 0.2% as the manager's stake triples, the impact on corporate policies is far more dramatic. The distortion in the investment rate (red line) increases by an order of magnitude, reaching nearly 2% relative to the baseline.

[Insert Figure 6 about here]

This divergence has important implications for the link between asset prices and real activity. The result implies that belief-driven distortions in investment are most severe precisely when firms are performing well and managerial are non rational. Consequently, managerial biases may act as a time-varying wedge that decouples investment from Tobin's q, particularly in mature, successful firms. This mechanism could worsen the observed disconnection between investment and valuation observed in the data, suggesting that the disconnection may be exacerbated by an equilibrium outcome of optimal contracting with biased agents (Gutiérrez and Philippon, 2017, Gashi, 2025).

[14] Another interpretation can be that the current dataset has limited coverage of the employees who are giving incentives compensation structure and this exercise can shed lights to the implication of this feature of the data.

7.5 External Validity - The Duke CFO Survey

A natural concern with the baseline estimation is that managerial guidance may be subject to strategic disclosure motives, potentially contaminating the estimation of belief parameters. To address this concern, I conduct a robustness exercise using an alternative source of managerial expectations: the Duke CFO Business Outlook Survey. This dataset provides a direct elicitation of CFO beliefs through confidential surveys, which are less susceptible to strategic manipulation than publicly disclosed guidance. The key advantage of this approach is that survey responses are not disclosed to the market.

The objective of this robustness exercise is twofold. First, I verify that the structural estimates are qualitatively consistent across the two independent data sources. Second, I assess whether the quantitative magnitudes are comparable, which would strengthen the external validity of the main findings. If the estimates from publicly disclosed guidance and confidential survey responses align, this would provide compelling evidence that the identified belief distortions reflect genuine managerial expectations.

To implement this analysis, I restrict the analysis to the subset of firms for which both managerial guidance (from IBES/FactSet) and CFO survey responses are available in the same fiscal year. This matched sample allows for a direct comparison.

The dataset of CFO expectations in the robustness section comes from the Duke CFO Business Outlook Survey, first led by John Graham and Campbell Harvey at Duke University's Fuqua School of Business (Graham et al., 2011 Ben-David et al., 2013). The survey is conducted quarterly and asks CFOs to report their prospects for future firm performance and policy, as well as their expectations for the future US economy and market performance.

The survey is conducted at Duke University until 2020Q1, after which it is in partnership with the Federal Reserve Banks of Richmond and Atlanta[15].

For the purpose of the paper, to match CFO expectations with firm fundamentals, I need the firm identifiers. While the survey does not require CFOs to identify themselves, some respondents voluntarily provide this information. I obtain this quarterly firm-level CFO survey data directly from John Graham, with privacy restrictions. Hence, the dataset used in the paper runs from 2001Q4 to 2020Q1 because it is only available before the collaboration with the Fed.

The first asks CFOs about their short-term (1-year) and long-term (10-year) S&P 500 return expectations. The original question is as follows:

> Over the next 12 month, I expect the annual S&P 500 return will be: _____%;
>
> Over the next 10 years, I expect the average annual S&P 500 return will be: _____%.

Second, the survey asks managers to report their prospects of the future 1-year growth of many corporate variables. The original question is as follows:

> Relative to the previous 12 months, what will be your company's PERCENTAGE CHANGE during the next 12 months? (e.g., +3%, 0%, -1%, etc.) [Leave blank if not applicable.]
>
> _____% Capital spending; _____% Earnings; _____% Research and development spending; _____% Revenue; _____% Cash on balance sheet;
>
> _____% Dividends; _____% Share repurchases[16].

[15] For details about the survey, visit https://www.richmondfed.org/research/national_economy/cfo_survey.

[16] Only selected items that are used in the analyses are listed. Check the survey website for a complete list of variables asked.

I match managers to firms using permno code to Compustat and IBES for firm-level analyses, so the sample is restricted to the set of identified firms. For the responses with non-missing permno code, I match the quarter forecast with the fiscal year that is forecasted. Moreover, I back out the implied expected sales given the answer in the survey. I keep only the firm-year for which I have both the CFO survey and the IBES Guidance to have a comparison. This results in 44 firms and 114 firm-year responses.

To assess the robustness of the baseline estimates, I re-estimate the structural model on the matched sample for which both managerial guidance and CFO survey responses are available. The estimation procedure follows the methodology described before. The key difference is that I now estimate the model twice for the same set of firms: once using publicly disclosed guidance as the measure of managerial expectations, and once using the confidential CFO survey responses. This design allows for a direct comparison of the structural estimates across the two data sources, holding the sample composition fixed.

The objective of this exercise is to verify whether the two independent measures of managerial expectations yield consistent inference about the belief parameters α (optimism) and β (confidence). If the estimates align qualitatively and quantitatively, this would provide strong external validation for the main findings, as it would indicate that the identified belief distortions are not driven by strategic disclosure motives inherent to managerial guidances.[17] I report in the figures below the results for the estimated parameters.

[17]It is worth highlighting that the CFO survey will allow to measure the bias on the first moment. For comparability, I compare this with the mid-point guidance. Therefore, the estimation of the bias on the second moment will be given by the dynamics of the state variable (scaled deferred compensation).

[Insert Figure 7,8, 9, and 10 about here]

Figures 7–10 present the structural parameter estimates obtained from the matched sample of firms with both managerial guidance and CFO survey data. Each figure compares the point estimate and 95% confidence interval from the two data sources: the publicly disclosed guidance (blue bars) and the confidential CFO survey responses (red bars).

The central finding is that the belief parameters exhibit strong qualitative agreement across the two independent data sources. Figure 7 reports the estimates of α, the optimism parameter. Both data sources yield positive point estimates: 0.033 from guidance and 0.012 from the CFO survey, indicating that managers are optimistic on average, consistent with the baseline results. Although the guidance estimate is somewhat larger in magnitude, both estimates are statistically distinguishable from zero and the directional agreement is unambiguous.

Figure 8 displays the corresponding estimates of β, the confidence parameter. Both data sources yield negative estimates: -0.005 from guidance and -0.023 from the CFO survey, confirming that managers are overconfident about the volatility of future outcomes. This directional agreement across two distinct measurement approaches strengthens the conclusion that overconfidence is a robust feature of managerial expectations, and is consistent with the systematic pattern of overconfidence documented in the main analysis as well as with the broader behavioral finance literature, which has long identified overconfidence as a pervasive feature of managerial decision-making.

Figures 9 and 10 report the estimates of the remaining structural parameters: the capital adjustment cost λ and the volatility σ. The CFO survey yields a higher estimate of λ (4.490 versus 2.637) and a lower estimate of σ (0.070 versus 0.120). Importantly, these parameters are not the primary object of interest for the robustness exercise; the focus is on the belief parameters α and β, which exhibit strong concordance.

Taken together, these results provide external validation for the main findings of the paper. The qualitative alignment of α and β across two independent data sources: one based on public disclosures subject to potential strategic motives, the other based on confidential survey responses, substantially mitigates concerns about measurement contamination.[18] The quantitative magnitudes are also broadly comparable, with both sources indicating optimism and moderate overconfidence. This robustness check therefore reinforces the credibility of the baseline estimates and supports the paper's central conclusion that managerial belief distortions are economically meaningful determinants of executive compensation and corporate investment.

8 Conclusion

This paper studies how managerial subjective beliefs affect incentives and firm policies. I develop a tractable dynamic contracting model in which the principal and the manager may disagree about future outcomes, and I characterize how the optimal contract adjusts to accommodate such beliefs. The analysis shows that managerial beliefs distort the perceived value of deferred compensation, thereby reshaping both the level and structure of executive pay. Through estimation on managerial forecast data, I document substantial cross-sectional variation in beliefs across firms and industries. Despite this heterogeneity, the optimal contract endogenously adjusts to account for the disagreement between managers and shareholders. The counterfactual analysis shows that eliminating belief distortions would marginally increase firm value, with limited effects on investment. These results suggest that managerial biases, while economically meaningful in shaping incentives, have limited aggregate impact as optimal contracts are in place. This helps rationalize the persistence of behavioural biases among top executives in equilibrium.

[18] It is possible to show that also the counterfactual exercises implies by these estimates lead to similar outcomes and magnitudes.

Future research can build on this framework along several dimensions. First, extending the model to incorporate learning about managerial skill would allow the principal to dynamically update beliefs enriching the analysis. Second, exploring beliefs at different maturities could shed light on the horizon of corporate policies (long- vs short-termism). Finally, quantifying the role of belief distortions and the cross-section of incentives for other economic agents (asset managers, bank executives...), would provide a broader understanding of how expectations propagate through financial markets and decision-making.

References

Adrian T. and Westerfield M. (2009). "Disagreement and Learning in a Dynamic Contracting Model". *The Review of Financial Studies* 22.10, pp. 3873–3906.

Ai H., Kiku D., and Li R. (2023). "A Quantitative Model of Dynamic Moral Hazard". *The Review of Financial Studies* 36.4, pp. 1408–1463.

Ai H. and Li R. (2015). "Investment and CEO Compensation under Limited Commitment". *Journal of Financial Economics* 116.3, pp. 452–472.

Barrero J. M. (2022). "The Micro and Macro of Managerial Beliefs". *Journal of Financial Economics* 143.2, pp. 640–667. DOI: 10.1016/j.jfineco.2021.06.007. URL: https://doi.org/10.1016/j.jfineco.2021.06.007.

Ben-David I., Graham J. R., and Harvey C. R. (2013). "Managerial Miscalibration". *Quarterly Journal of Economics* 128.3, pp. 1547–1584. URL: https://doi.org/10.1093/qje/qjt024.

Biais B., Mariotti T., Plantin G., and Rochet J.-C. (2007). "Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications". *The Review of Economic Studies* 74.2, pp. 345–390.

Biais B., Mariotti T., Rochet J.-C., and Villeneuve S. (2010). "Large Risks, Limited Liability, and Dynamic Moral Hazard". *Econometrica* 78.1, pp. 73–118.

Bolton P., Chen H., and Wang N. (2011). "A Unified Theory of Tobin's q, Corporate Investment, Financing, and Risk Management". *The Journal of Finance* 66.5, pp. 1545–1578.

Bordalo P., Gennaioli N., La Porta R., and Shleifer A. (2017). "Diagnostic Expectations and Stock Returns". *NBER Working Paper Series* 23863. URL: http://www.nber.org/papers/w23863.

Bordalo P., Gennaioli N., and Shleifer A. (2018). "Diagnostic Expectations and Credit Cycles". *Journal of Finance* 73.1, pp. 199–227. DOI: 10.1111/jofi.12586. URL: https://doi.org/10.1111/jofi.12586.

Bouchaud J.-P., Krüger P., Landier A., and Thesmar D. (2019). "Sticky Expectations and the Profitability Anomaly". *Journal of Finance* 74.2, pp. 639–674. DOI: 10.1111/jofi.12734. URL: https://doi.org/10.1111/jofi.12734.

Bustamante M. C., Cujean J., and Frésard L. (2022). "Knowledge Cycles and Corporate Investment". *Working Paper / SSRN*. URL: https://ssrn.com/abstract=3418171.

Coibion O., Gorodnichenko Y., and Kumar S. (2018). "How Do Firms Form Their Expectations? New Survey Evidence". *American Economic Review* 108.9, pp. 2671–2713. URL: https://doi.org/10.1257/aer.20151299.

Cooper I. (2006). "Asset Pricing Implications of Nonconvex Adjustment Costs and Irreversibility of Investment". *The Journal of Finance* 61.1, pp. 139–170.

De La O R. and Myers S. (2021). "Subjective Cash Flow and Discount Rate Expectations". *Journal of Finance* 76.3, pp. 1339–1387. DOI: 10.1111/jofi.13016. URL: https://doi.org/10.1111/jofi.13016.

– (2024). "Which Subjective Expectations Explain Asset Prices?" *Review of Financial Studies* 37.6, pp. 1929–1978. DOI: 10.1093/rfs/hhae009. URL: https://doi.org/10.1093/rfs/hhae009.

De Marzo P. and Sannikov Y. (2017). "Learning, Termination, and Payout Policy in Dynamic Incentive Contracts". *Review of Economic Studies*.

DeMarzo P. M., Fishman M. J., He Z., and Wang N. (2012). "Dynamic Agency and the Q Theory of Investment". *The Journal of Finance* 67.6, pp. 2295–2349.

DeMarzo P. M. and Sannikov Y. (2006). "Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model". *The Journal of Finance* 61.6, pp. 2681–2724.

Dittmann I. and Maug E. (2007). "Lower Salaries and No Options? On the Optimal Structure of Executive Pay". *Journal of Finance* 62.1, pp. 303–343. DOI: 10.1111/j.1540-6261.2007.01208.x. URL: https://doi.org/10.1111/j.1540-6261.2007.01208.x.

Dittmann I., Maug E., and Spalt O. (2010). "Sticks or Carrots? Optimal CEO Compensation when Managers Are Loss Averse". *Journal of Finance* 65.6, pp. 2015–2050. DOI: 10.1111/j.1540-6261.2010.01609.x. URL: https://doi.org/10.1111/j.1540-6261.2010.01609.x.

Epstein L. and Ji S. (2013). "Ambiguous Volatility and Asset Pricing in Continuous Time". *Review of Financial Studies*.

Frésard L., Mancini L., Schroth E., and Sinno D. (2023). "How Do Firms Choose Between Growth and Efficiency?" *Working Paper*.

Gao F., Zighuo H., Wei B., and Yu J. (2017). "Optimal Long-term Contracting with Learning". *The Review of Financial Studies* 30, pp. 2006–2065.

Gashi D. (2025). "Subjective Beliefs and the Q-Theory of Investments". *Working Paper*.

Gayle G.-L. and Miller R. A. (2009). "Has Moral Hazard Become a More Important Factor in Managerial Compensation?" *American Economic Review* 99.5, pp. 1740–1769.

Gervais S., Heaton J. B., and Odean T. (2011). "Overconfidence, Compensation Contracts, and Capital Budgeting". *The Journal of Finance* 66, pp. 1735–1777.

Graham J. R., Harvey C. R., and Puri M. (2013). "Managerial Attitudes and Corporate Actions". *Journal of Financial Economics* 109.1, pp. 103–121. DOI: 10.1016/j.jfineco.2013.01.010. URL: https://EconPapers.repec.org/RePEc:eee:jfinec:v:109:y:2013:i:1:p:103-121.

Graham J. R., Harvey C. R., and Puri M. (2015). "Capital Allocation and Delegation of Decision-Making Authority within Firms". *Journal of Financial Economics* 115.3, pp. 449–470. DOI: 10.1016/j.jfineco.2014.10.011. URL: https://doi.org/10.1016/j.jfineco.2014.10.011.

Graham J. R., Li S., and Qiu J. (2011). "Managerial Attributes and Executive Compensation". *The Review of Financial Studies*. URL: https://ssrn.com/abstract=1267189.

Greenwood R. and Hanson S. G. (2015). "Waves in Ship Prices and Investment". *The Quarterly Journal of Economics* 130.1, pp. 55–110. URL: https://doi.org/10.1093/qje/qju035.

Greenwood R. and Shleifer A. (2014). "Expectations of Returns and Expected Returns". *Review of Financial Studies* 27.3, pp. 714–746. URL: https://doi.org/10.1093/rfs/hht082.

Gryglewicz S. and Hartman-Glaser B. (2020). "Investment Timing and Incentive Costs". *The Review of Financial Studies* 33.1, pp. 309–357. DOI: 10.1093/rfs/hhz051.

Gryglewicz S., Mancini L., Morellec E., Schroth E., and Valta P. (2021). "Understanding Cash Flow Risk". *The Review of Financial Studies* 35.8, pp. 3922–3972. DOI: 10.1093/rfs/hhab127. URL: https://doi.org/10.1093/rfs/hhab127.

Guiso L. and Parigi G. (1999). "Investment and Demand Uncertainty". *Quarterly Journal of Economics* 114.1, pp. 185–227. DOI: 10.1162/003355399555981. URL: https://doi.org/10.1162/003355399555981.

Gutiérrez G. and Philippon T. (2017). "Declining Competition and Investment in the U.S." Working Paper Series 23583. DOI: 10.3386/w23583. URL: http://www.nber.org/papers/w23583.

Hackbarth D. and Johnson T. (2015). "Real Options and Risk Dynamics". *The Review of Economic Studies* 82.4, pp. 1449–1482.

He Z. (2009). "Optimal Executive Compensation when Firm Size Follows Geometric Brownian Motion". *The Review of Financial Studies* 22.3, pp. 859–892.

– (2011). "A Model of Dynamic Compensation and Capital Structure". Journal of Financial Economics.

Ma Y., Ropele T., Sraer D., and Thesmar D. (2024). "A Quantitative Analysis of Distortions in Managerial Forecasts". *Working Paper.*

Malenko A. (2019). "Optimal Dynamic Capital Budgeting". *The Review of Economic Studies* 86.4, pp. 1747–1778.

Malmendier U. and Tate G. (2015). "Behavioral CEOs: The Role of Managerial Overconfidence". *Journal of Economic Perspectives* 29.4, pp. 37–60. URL: https://doi.org/10.1257/jep.29.4.37.

Miao J. and Rivera A. (2016). "Robust Contracts in Continuous Time". *Econometrica* 84.4, pp. 1405–1440.

Nikolov B., Schmid L., and Steri R. (2021). "The Sources of Financing Constraints". *Journal of Financial Economics* 139.2, pp. 478–501.

Nikolov B. and Whited T. M. (2014). "Agency Conflicts and Cash: Estimates from a Dynamic Model". *The Journal of Finance* 69, pp. 1883–1921.

Otto C. (2014). "CEO Optimism and Incentive Compensation". *Journal of Financial Economics.*

Otto C. A. (2014). "CEO Optimism and Incentive Compensation". *Journal of Financial Economics* 114.2, pp. 366–404. URL: https://doi.org/10.1016/j.jfineco.2014.06.006.

Page B. (2018). "CEO Attributes, Compensation, and Firm Value: Evidence from a Structural Estimation". *Journal of Financial Economics* 128, pp. 378–401.

Peters R. and Taylor L. (2017). "Intangible capital and the investment-q relation". *Journal of Financial Economics.*

Rivera A. (2023). "Contracting with a Present-Biased Agent: Sannikov Meets Laibson". *Working Paper.*

Sannikov Y. (2008). "A Continuous-Time Version of the Principal-Agent Problem". *The Review of Economic Studies* 75.3, pp. 957–984.

Stein J. C. (1996). "Rational Capital Budgeting in an Irrational World". *Journal of Business* 69.4, pp. 429–455. DOI: 10.1086/209699. URL: https://doi.org/10.1086/209699.

Tanaka M., Bloom N., David J. M., and Koga M. (2020). “Firm Performance and Macro Forecast Accuracy”. *Journal of Monetary Economics* 114.C, pp. 26–41. DOI: 10.1016/j.jmoneco.2019.02.008. URL: https://doi.org/10.1016/j.jmoneco.2019.02.008.

Taylor L. A. (2013). “CEO Wage Dynamics: Estimates from a Learning Model”. *Journal of Financial Economics* 108.1, pp. 79–98.

Ward C. (2025). “Agency in Intangibles”. – –.–, pp. –.

Figure 1: Comparative Statics. This figure displays the model's comparative statics with respect to the belief parameter α. The two panels plot, as functions of the scaled continuation utility w: firm value $p(w)$ (left) and the investment rate $i(w)$ (right). Positive (negative) values of α correspond to optimistic (pessimistic) managerial beliefs.

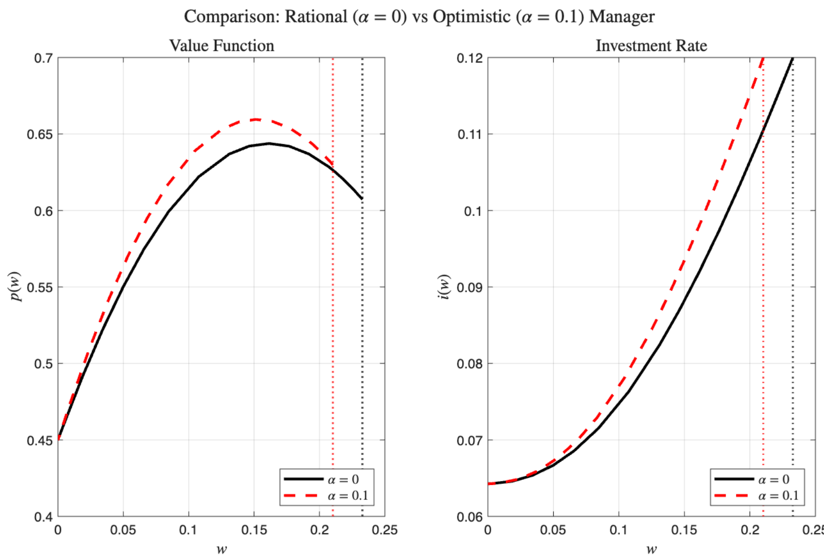

Figure 2: Cost of Belief Distortions. This figure plots the percentage change in firm value relative to the rational benchmark ($\alpha = 0$) for varying degrees of managerial optimism ($\alpha > 0$) and pessimism ($\alpha < 0$). The vertical axis measures the percent deviation $\Delta p(w)/p(w)|_{\alpha=0}$.

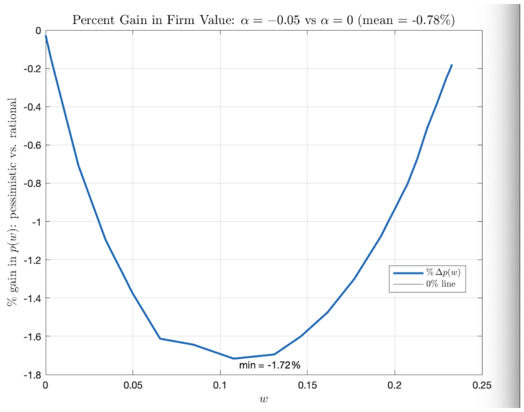

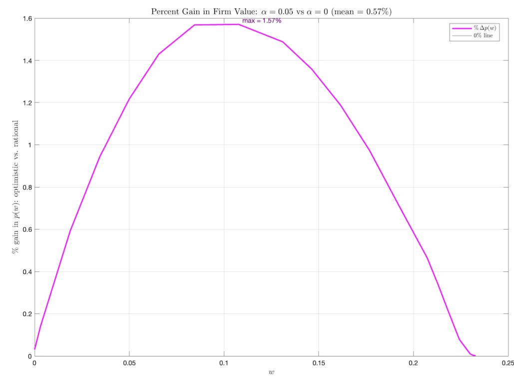

Figure 3: Alpha by industry. The figure reports the average estimated α (optimism/pessimism) by industry, with 95% confidence intervals. Positive values indicate optimism ($\alpha > 0$), while negative values indicate pessimism ($\alpha < 0$).

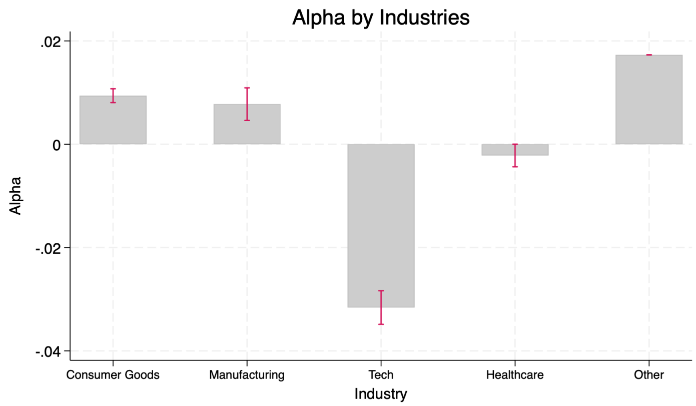

Figure 4: Beta by industry. The figure reports the average estimated β (over-/underconfidence) by industry, with 95% confidence intervals. Negative values indicate overconfidence ($\beta < 0$), while positive values indicate underconfidence ($\beta > 0$).

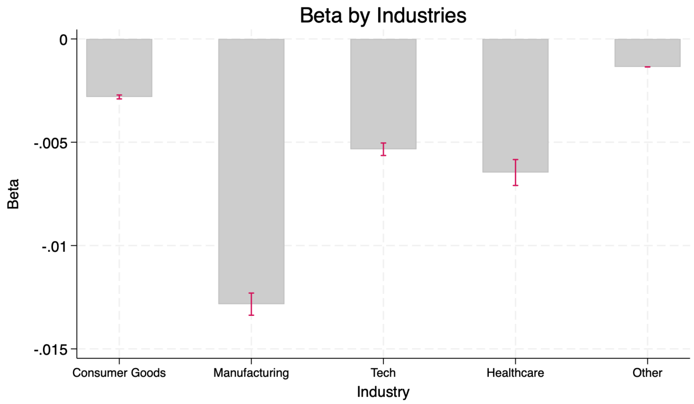

Figure 5: Managerial incentives implied by subjective beliefs across industries. The figure reports the belief-adjusted incentive intensity $\lambda \propto (1+\beta)/(1+\alpha)$ *by industry. Values above one indicate that stronger incentives are required relative to the rational benchmark.*

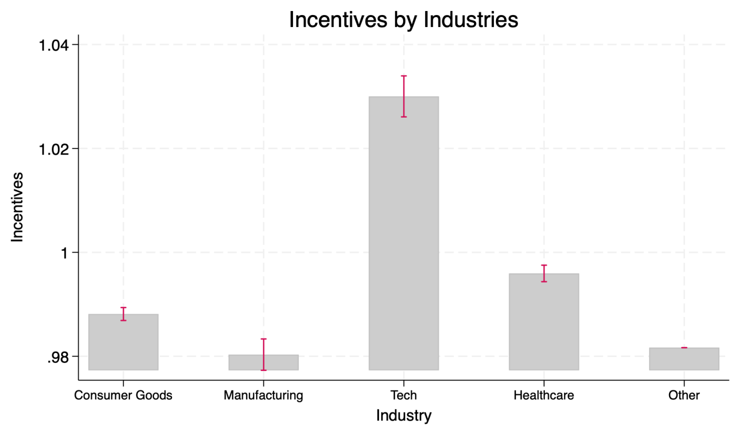

Figure 6: Counterfactual Impact vs. Scaled State Variable. This figure illustrates the state-dependency of belief distortions. It plots the percentage change in Firm Value (blue line, left axis) and Investment Rate (red line, right axis) that would result from replacing a biased manager with a rational one, as a function of the state variable multiplier ξ. The horizontal axis represents ξ, a scalar multiplier applied to the estimated state variable w (scaled continuation utility), simulating scenarios where the manager's claim on the firm increases relative to the capital stock (e.g., following a sequence of positive shocks). $\xi = 1$ corresponds to the sample average.

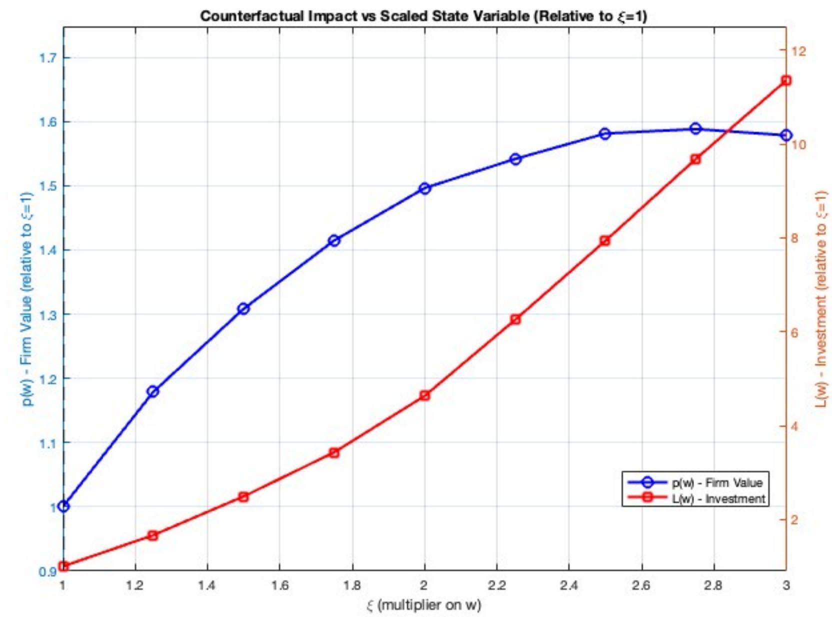

Figure 7: Estimated α Parameter: Guidance Sample vs. CFO Survey Sample. This figure compares the distribution of estimated parameters α across the management guidance sample (baseline) and the Duke CFO Survey sample.

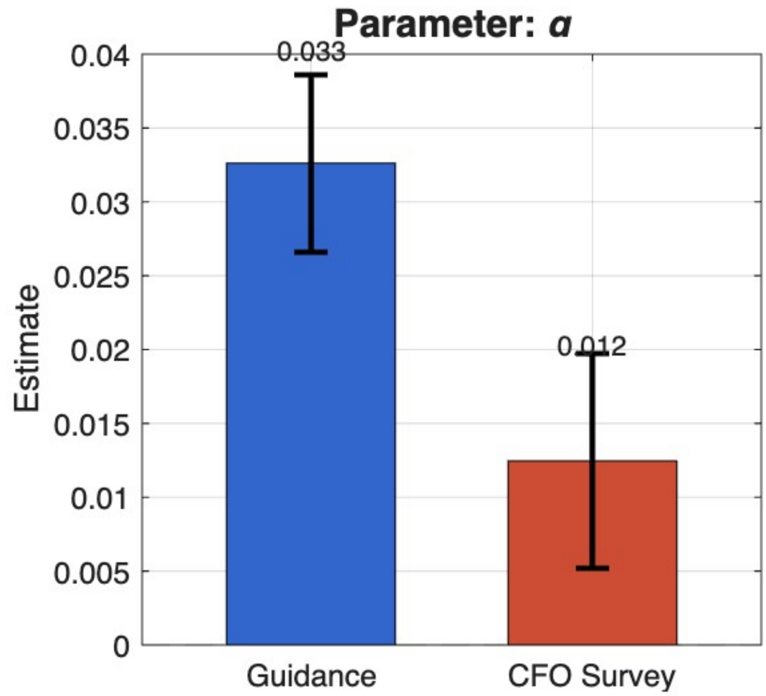

Figure 8: Estimated parameter β: Guidance sample and CFO Survey Sample. This figure compares the distribution of estimated parameters β across the management guidance sample (baseline) and the Duke CFO Survey sample.

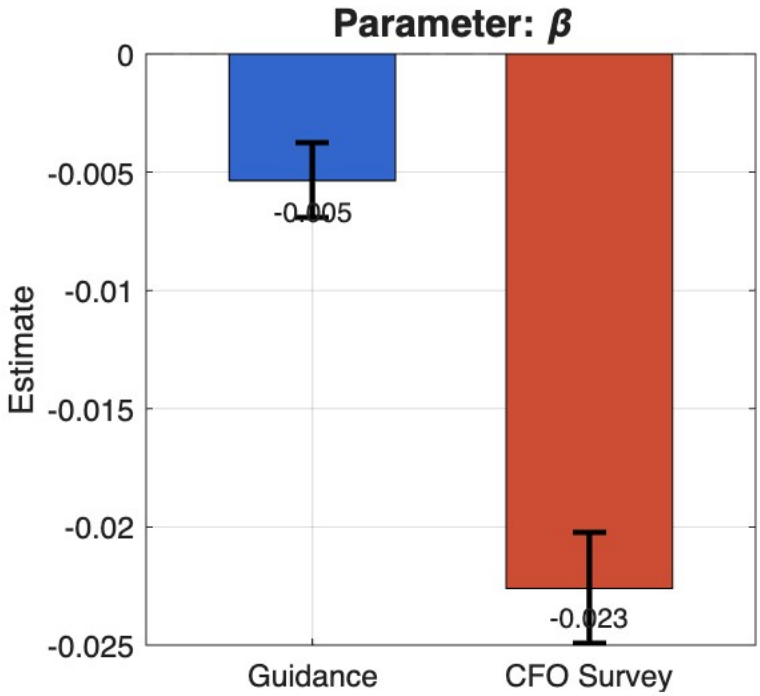

Figure 9: Estimated parameter λ: Guidance sample and CFO Survey Sample. This figure compares the distribution of estimated incentive parameters λ across the management guidance sample (baseline) and the Duke CFO Survey sample.

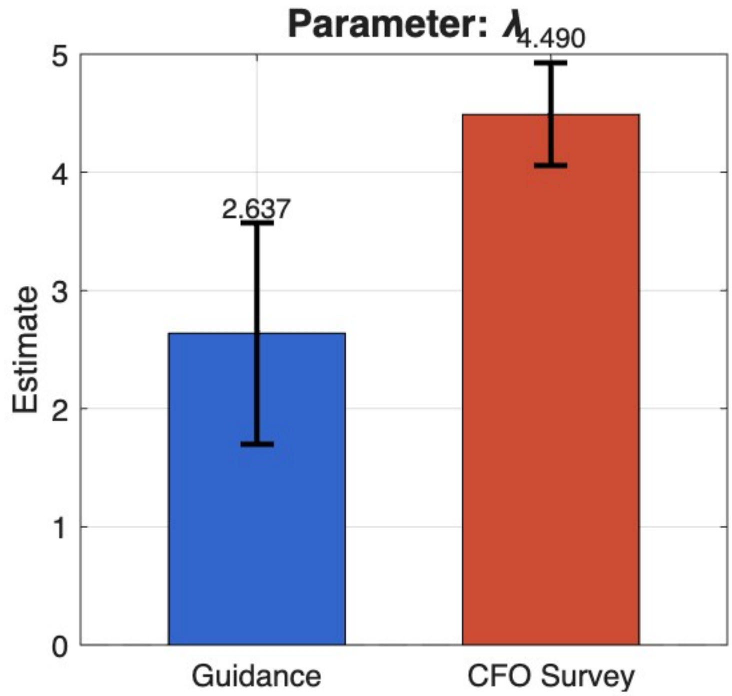

Figure 10: Estimated parameter σ: Guidance sample and CFO Survey Sample. This figure compares the distribution of estimated volatility parameters σ across the management guidance sample (baseline) and the Duke CFO Survey sample.

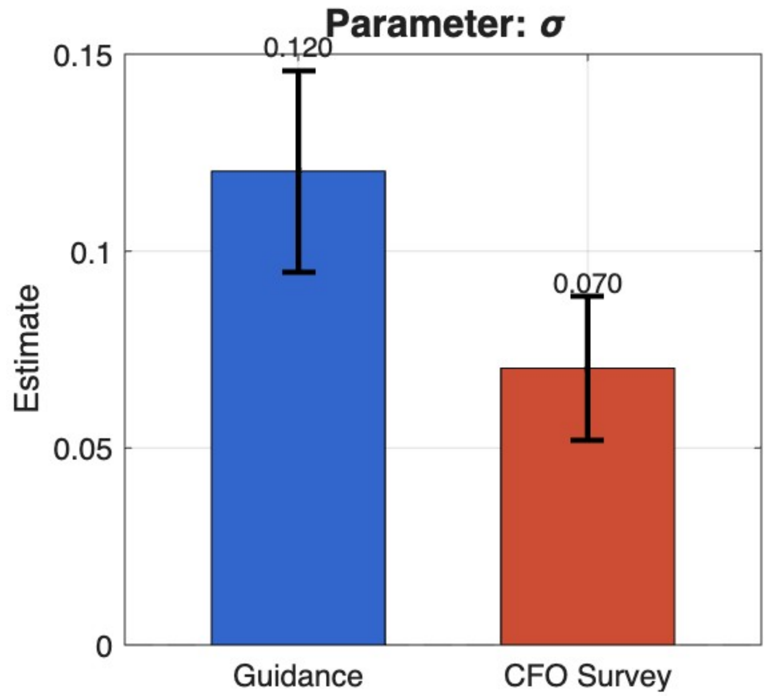

Table 1: Summary Statistics. This table presents descriptive statistics of the main variables used in the analysis. The descriptive statistics are: mean; standard deviation; and the percentiles p25, p50, and p75. The sample covers the period 2005 to 2019.

	Mean	Std	p25	p50	p75
q	1.19	1.26	0.44	0.86	1.55
$TotalCompensation$	11.80	10.38	5.15	9.41	15.21
$Deferred\ Compensation$	7.03	7.47	2.40	5.24	9.47
$investmentRate$	0.19	0.10	0.12	0.17	0.24
$forecastedSalesLow$	1.21	0.74	0.70	1.02	1.52
$forecastedSalesHigh$	1.5	0.76	0.72	1.04	1.56
$Sales/Turnover$	1.23	0.75	0.70	1.04	1.53
$Assets$ (Total)	1990.03	4472.61	321.30	652.61	1782.44

Table 2: Sales Forecasts Predicting Realized Sales. This table reports estimates from the regression $Sales_{i,t} = \alpha + \beta F_{i,t-1}\left[Sales_{i,t}\right] + \epsilon_{i,t}$. *All variables are scaled by total assets. Year fixed effects are included in both specifications; column (2) additionally includes industry fixed effects. Standard errors (in parentheses) are clustered by year in column (1) and two-way clustered by industry and year in column (2). ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.*

	(1)	(2)
	$Sales_{i,t}$	
$F_{i,t-1}[Sales_{i,t}]$	0.91***	0.91***
	(0.02)	(0.02)
Industry FE	N	Y
Year FE	Y	Y
Adj. R^2	0.96	0.97
N	1,704	1,704

*Table 3: Managerial Beliefs and Incentive Provision. This table reports estimates of the regression $Incentives_{i,t} = \alpha + \beta_1 FE_{i,t} + \beta_2 ExcessVol_{i,t} + \epsilon_{i,t}$. $ExcessMean$ denotes the forecast error (optimism), and $ExcessVol$ the excess variance of forecast errors (underconfidence). All variables are scaled by assets. Standard errors (in parentheses) are double-clustered by firm and year. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.*

	Dependent variable: Deferred Compensation (Incentives)			
	(1)	(2)	(3)	(4)
$ExcessSubjectiveMean$	-0.0218***		-0.023***	-0.0157**
	(0.0039)		(0.0044)	(0.0044)
$ExcessSubjectiveVol$		0.0109***	0.0134***	0.01796**
		(0.0031)	(0.0031)	(0.006)
Controls	N	N	N	Y
Industry FE	N	N	N	Y
Year FE	N	N	N	Y
Adj. R^2	0.015	0.01	0.02	0.30
N	1,704	1,704	1,704	1,704

Table 4: Estimation results for the groups. This table reports the estimated parameters (mean), their standard deviation and the percentiles p25, p50 and p75.

	Mean	Std. Dev.	p25	p50	p75
λ_k	1.671	0.532	1.399	1.776	1.920
σ	0.084	0.014	0.073	0.078	0.087
α	-0.009	0.039	-0.015	0	0.009
β	-0.006	0.005	-0.006	-0.004	-0.002

Table 5: Standard errors of the estimated parameters. This table reports the standard errors (mean), their standard deviation and the percentiles p25, p50 and p75.

	Mean	Std. Dev.	p25	p50	p75
se-λ	0.0225	0.0188	0.0069	0.0186	0.0432
se-σ	0.0003	0.0004	0.0001	0.0001	0.0013
se-α	0.00131	0.00084	0.00060	0.00131	0.00205
se-β	0.00044	0.00040	0.00020	0.00028	0.00051

Table 6: Likelihood ratio tests for belief parameters. This table reports likelihood ratio tests comparing the unrestricted model against nested specifications. Critical values correspond to the 5% significance level of the χ^2 distribution with the appropriate degrees of freedom.

Null Hypothesis	Rejection Rate
$\alpha = 0$	57.9%
$\beta = 0$	62.8%
$\alpha = \beta = 0$	74.4%

Table 7: This table reports the average incentive distortions ratio $\lambda^{irr} = (1+\beta)/(1+\alpha)$ *for firms below and above the mean of each firm characteristic. A value above (below) one indicates that the manager is paid more (less) than its rational counterpart. Age is the number of years since IPO; Size is the log of real total assets; Intangibility is intangible capital over total capital; Fin. Constraint is the financing delay index of Hoberg and Maximovic; Inst. Ownership is the share of institutional investors. Statistical significance based on a two-sample t-test:* $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

λ^{irr}	Below Mean	Above Mean	Difference
Age	1.0076	1.0012	-0.0064***
Size (log assets)	1.0069	1.0027	-0.0042**
Intangibility	0.9967	1.0101	0.0133***
Fin. Constraint	0.9996	1.0069	0.0073***
Inst. Ownership	1.0119	1.0003	-0.0116***

Table 8: Counterfactual Analysis: Heterogeneity in Firm Value Impact. This table reports the average counterfactual changes for groups with the lowest and highest percentage change in firm value $p(w)$ when replacing the biased manager with a rational one ($\alpha = \beta = 0$). The column "Lowest $\Delta p(w)$" corresponds to firms where the biased manager generates higher value than a rational one; "Highest $\Delta p(w)$" corresponds to firms where the biased manager destroys value. $\Delta p(w)$ and $\Delta L(w)$ denote the percentage change in firm value and investment, respectively. α captures optimism (positive) or pessimism (negative), β captures underconfidence (positive) or overconfidence (negative), and $(1+\alpha)/(1+\beta)$ is the effective incentive wedge (rational benchmark = 1).

	Lowest $\Delta p(w)$ *(Benefit of Bias)*	Highest $\Delta p(w)$ *(Cost of Bias)*
$\Delta p(w)$ (%)	-0.15	0.06
$\Delta L(w)$ (%)	-0.03	0.02
α	0.02	-0.02
β	-0.01	-0.01
Bias Ratio: $(1+\beta)/(1+\alpha)$	1.03	0.99

Appendices

A Theoretical Analysis

A.1 Regularity Conditions

To ensure that the problem is well-behaved I impose the following regularity conditions. Square integrability of the payout process $\{C\}$:

$$E\left[\left(\int_0^{\tau} e^{-\gamma s} dC_s\right)^2\right] < \infty.$$

$$E\left[\int_0^{T} \left(e^{-rt} A_t K_t\right)^2 dt\right] < \infty \quad \text{for all } T > 0,$$

$$\lim_{T \to \infty} E\left(e^{-rT} A_T K_T\right) = 0.$$

Therefore, the process i_t is of bounded variation.

A.2 Continuation Payoff

We operate in the probability measure of the manager and drop the superscript M, we therefore derive dW_t (as usual):

- Define V_t as the value process of the agent's discounted wages from time 0:

$$V_t \equiv E_t\left[\int_0^\tau e^{-\gamma s} dC_s\right]$$

  V_t is a martingale.

- By the Martingale Representation Theorem:

$$V_t = V_0 + \int_0^t e^{-\gamma s} K_t A_t \psi_s dZ_s$$

- Assuming the agent works ($e_s = 1$), $dZ_s = \frac{dA_s - \mu A_s ds}{\sigma_a A_s}$. Substituting:

$$\begin{aligned} V_t &= V_0 + \int_0^t e^{-\gamma s} K_s A_s \psi_s \left(\frac{dA_s - \mu A_s ds}{\sigma_a A_s}\right) \\ &= V_0 + \int_0^t e^{-\gamma s} K_s \frac{\psi_s}{\sigma_a}(dA_s - \mu A_s ds) \end{aligned}$$

  In differential form:

$$dV_t = e^{-\gamma t} K_t \frac{\psi_t}{\sigma_a}(dA_t - \mu A_t dt) \tag{27}$$

- Therefore, we get as usual:

$$dW_t = \gamma W_t dt - dC_t + \frac{\psi_t}{\sigma_a} K_t (dA_t - \mu A_t dt) \tag{28}$$

  and this is continuation payoff from the point of view of the agent

- with similar math, the **Manager Perceived Continuation Payoff** $\tilde{W}$ evolves:

$$dW_t^M = \gamma W_t^M dt - dC_t + \frac{\psi_t}{\sigma_a(1+\beta)} K_t(dA_t - (1+\alpha)\mu A_t dt) \tag{29}$$

A.3 Model Solution

Optimality principle + ITO lemma and the HJB in $P(W, A, K) = P$ (drop t for now) becomes:

$$\begin{aligned} rP = {} & (\mu - c(i))\, AK + P_A\, \mu A + P_K\, (i - \delta_k)K \\ & + P_W\, \gamma W + \tfrac{1}{2}\, P_{AA}\, \sigma_A^2 A^2 \\ & + \tfrac{1}{2}\, P_{WW}\, (\sigma_a\, \lambda^{\text{Biased}}\, AK)^2 + P_{AW}\left(\sigma_a^2\, \lambda^{\text{Biased}}\, A^2 K\right) \end{aligned} \tag{30}$$

Use homogeneity and $P(A, K, W) = AKp\left(\frac{W}{AK}\right) = AKp(w)$ where I defined: $w = \frac{W}{AK}$, remember when solving that $W = wAK$.

- $P_W = p'(w)$
- $P_A = K[p(w) - wp'(w)]$
- $P_K = A[p(w) - wp'(w)]$
- $P_{WW} = p''(w)\frac{1}{AK}$
- $P_{AA} = \frac{Kw^2}{A}p''(w)$
- $P_{AW} = -\frac{w}{A}p''(w)$

Solve and substitute and the final HJB in x is:

$$rp(w) = [\mu - c(i)] + p(w)[\mu + i - \delta_k] + wp'(w)[\gamma - (i - \delta_k + \mu)] + \frac{1}{2}p''(w)\sigma_a^2(w - \lambda^{Biased})^2 \quad (31)$$

when $w \in [0, \bar{w}]$ and $dC = 0$. Moreover, I have the usual three conditions:

- $p'(\overline{w}) = -1$
- $p''(\overline{w}) = 0$
- $p(0) = \ell$ (Liquidation)

A.4 Verification Theorem and Concavity

Proof for Verification

I verify the optimality of the contract Π^* among all contracts Π satisfying incentive compatibility. To do so, I show that the principal obtains under any contract $\Pi \in IC$ at most (scaled) payoff $\hat{G}(\Pi)/(A_0K_0) \leq p(w^*)$, with equality if and only if $\Pi = \Pi^*$. Here, $p(\cdot)$ solves the HJB with corresponding payout threshold $\bar{w}$ and $w_0 = w^*$. Consider any incentive-compatible contract $\Pi \equiv (\{C\}, \{i\}, \tau)$. For any $t \leq \tau$, define its auxiliary gain process G as

$$G_t(\Pi) = \int_0^t e^{-ru}\big(dY_u - c(i_u)A_uK_u\,du\big) - \int_0^t e^{-ru}dC_u + e^{-rt}P(W_t, A_t, K_t),$$

where the agent's continuation payoff evolves according to Eq. (7). Recall that $w_t = \frac{W_t}{A_t K_t}$, $P(W_t, A_t, K_t) = A_t K_t p(w_t)$, and $dc_t \equiv dC_t/(A_t K_t)$. Ito's lemma implies that for $t \le \tau$,

$$\begin{aligned}\frac{e^{rt}}{A_t K_t} dG_t(\Pi) = \Bigg[&- rp(w_t) + \mu - c(i_t) + (\mu + i_t - \delta_k)\big(p(w_t) - w_t p'(w_t)\big) \\ &+ \gamma w_t p'(w_t) + \frac{1}{2}\sigma_a^2(\lambda^{bias} - w_t)^2 p''(w_t)\Bigg] dt \\ &- (1 + p'(w_t))dc_t \\ &+ \sigma_a\big[1 + p(w_t) + p'(w_t)(\lambda^{bias} - w_t)\big] dZ_t.\end{aligned}$$

Under the optimal investment and incentives, the drift term in brackets equals zero. Other admissible policies make this term non-positive by the HJB (owing to the concavity result). The payout term is also non-positive since $p'(w_t) \ge -1$, and the expected value of the stochastic integral is zero under the regularity conditions. Therefore $\{G_t\}_{t \ge 0}$ follows a supermartingale, and it is a martingale under the optimal contract.

Furthermore, under Π, investors' expected payoff is

$$\tilde{G}(\Pi) \equiv E\left[\int_0^\tau e^{-ru}\big(dY_u - c(i_u)A_u K_u \, du\big) - \int_0^\tau e^{-ru} dC_u + e^{-r\tau} R A_\tau K_\tau\right].$$

As a result, investors' payoff is:

$$\begin{aligned}\tilde{G}(\Pi) &= E[G_\tau(\Pi)] \\ &= E\left[G_{t\wedge\tau}(\Pi) + 1_{\{t \le \tau\}}\left(\int_t^\tau e^{-rs}\big(dY_s - dC_s - c(i_s)A_s K_s \, ds\big) + e^{-r\tau} R A_\tau K_\tau - e^{-rt} P(W_t, A_t, K_t)\right)\right] \\ &= E[G_{t\wedge\tau}(\Pi)] \\ &\quad + e^{-rt} E\left[1_{\{t \le \tau\}} E_t\left(\int_t^\tau e^{-r(s-t)}\big(dY_s - dC_s - c(i_s)A_s K_s \, ds\big) + e^{-r(\tau - t)} R A_\tau K_\tau - P(W_t, A_t, K_t)\right)\right] \\ &\le G_0 + e^{-rt}(p^{FB} - R) E[A_t K_t].\end{aligned}$$

where p^{FB} is the scaled first-best value function. The last inequality follows from the super-martingale property of G_t, the fact that firm value under agency is bounded above by the first-best value, and the regularity conditions. Letting $t \to \infty$ gives

$$\tilde{G}(\Pi) \leq G_0 = P(W_0, A_0, K_0) = A_0 K_0 p(w_0).$$

Under the optimal contract, the inequality holds with equality. Since the principal chooses $w_0 = w^*$, where $p'(w^*) = 0$, this proves that Π^* is optimal.

Proof for Concavity

The proof is done by contradiction. Let $q(w) = p(w) - wp'(w)$. Using the optimal investment rule $i(w) = q(w)/\lambda_k$, the HJB becomes: $(r + \delta_k - \mu)p(w) = \mu + \frac{1}{2\lambda_k} q(w)^2 + wp'(w)[\gamma + \delta_k - \mu] + \frac{1}{2}\sigma_a^2 p''(w)(w - \lambda)^2$. Differentiating with respect to w and evaluating at any point where $p''(w) = 0$ gives $\frac{1}{2}\sigma_a^2 p'''(w)(w-\lambda)^2 = (r-\gamma)p'(w)$. Recall that $p'(\overline{w}) = -1$ and $p''(\overline{w}) = 0$. Therefore, at the payout boundary, $\frac{1}{2}\sigma_a^2 p'''(\overline{w})(\overline{w} - \lambda)^2 = \gamma - r > 0$, which implies $p''(\overline{w} - \epsilon) < 0$ for sufficiently small $\epsilon > 0$. Note that any point $w_1 \in (0, \overline{w}]$ satisfying $p''(w_1) = 0$ must also satisfy $p'(w_1) < 0$; otherwise, the HJB evaluated locally around w_1 would imply that the agency value reaches at least the first-best value, contradicting incentive costs and inefficient liquidation. Suppose now that p'' is not strictly negative on $[0, \overline{w})$ and define $w_1 = \sup\{w < \overline{w} : p''(w) \geq 0\}$. Then $p''(w_1) = 0$. Since $p'(w_1) < 0$, re-evaluating the differentiated HJB at w_1 gives $p'''(w_1) = \frac{2(r-\gamma)p'(w_1)}{\sigma_a^2 (w_1 - \lambda)^2} > 0$. Hence $p''(w_1 + \epsilon) > 0$ for sufficiently small $\epsilon > 0$, contradicting the definition of w_1. Therefore, $p''(w) < 0$ for all $w \in [0, \overline{w})$, so $p(w)$ is strictly concave on $[0, \overline{w}]$.

B Dataset construction

B.1 Forecasts Data

As explained in the main body of the paper, a key aspect of bringing the model to the data revolves around having accurate data about managerial and investors' expectations.

For the analysts I take the consensus from IBES and Factset. I consider only forecasts done in the short-term (12 months before the fiscal year). In there are more than one, I have the average. In a similar spirit, for managerial forecasts, I consider only short-term guidance. This is particularly important in this case. Indeed, managerial guidances are voluntarily disclosed (see Figure below) and might be manipulated by strategic communication with the public.

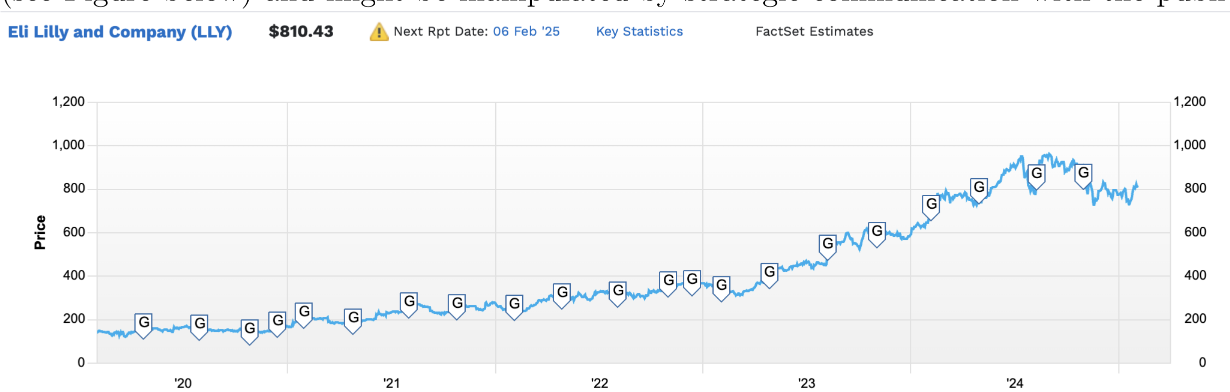

These values are expected to be between what the manager believes and what they want the public to believe. Nevertheless, in the short-term these concerns are mitigated as they will face immediate consequences. I want to highlight that I complement the common dataset used IBES with Factset to increase the panel size.

The two most commonly forecasted variables are Sales and EPS, with sales having the better coverage in my setup. Moreover, the choice of using sales is also motivated by not having to adjust the denominator of EPS, which could be problematic.

B.2 Data Cleaning

In this section, I will explain how I constructed the variables in the summary statistics. The variable cashflow is the sum of income before extraordinary items (ib), depreciation expenses (dp) and intangible investment(xrd+ 0.3*sga). The stock of physical capital is just PPEGT. The stock of intangible capital. In the paper, I will use "total capital", which is the sum of physical and intangible capital. The numerator of the definition of Tobin's Q is computed as usual ((prcc * csho) + (dlc+dltt) - act), while the denominator is Total capital for Total Q. I will rely on the data provided by Peters and L. Taylor, 2017 on WRDS. Lastly, size is the natural logarithm of total assets (atq).

C Model Estimation with Unscented Kalman Filter

This section provides a detailed exposition of the estimation method used in Section 3.2. We describe the state space model, the unscented Kalman filter to compute the likelihood function, and how we handle missing observations.

C.1 The state space model

The state space model in (12)–(18) consists of a transition equation and a measurement equation. The transition equation describes the discrete-time dynamics of the latent state process, which is the unobserved capital and labor stocks providing services for production. The measurement equation describes the relation between the state process and the observed data (earnings, capital, labor, investment, hiring) of firms that share the same state process in each group. To facilitate exposition, we use a standard notation in state space models, and present the model as if missing observations were absent (Appendix B.3 discusses how we handle missing observations).

The transition equation describes the discrete-time dynamic of the two-dimensional state process

$$x_t = [W_t,\, Z_t]',$$

with $'$ denoting transposition:

$$x_{t+1} = \phi_0 + \phi_1 x_t + w_t, \tag{32}$$

where $\phi_0 = [\mu_K \; \mu_L]'$, ϕ_1 is the identity matrix, $w_t \sim \mathcal{N}(0, Q)$, and Q is a diagonal covariance matrix with entries σ_K^2 and σ_L^2.

The measurement equation links the observed data to the state process and is given by

$$z_t = h(x_t) + v_t, \tag{33}$$

where the measurement error $v_t \sim \mathcal{N}(0, R)$. We consider groups of $N = 10$ firms and for each firm we obtain five variables, i.e. operating earnings, capital, labor, investment and hiring. The fifty-dimensional vector z_t collects all observed variables in every year t.

We allow measurement errors on each variable to have their own variance $\sigma^2_{v,1}, \ldots, \sigma^2_{v,5}$, resulting in a block diagonal covariance matrix R. Denoting by $x_{1,t} = \log(K_t)$ and $x_{2,t} = \log(L_t)$ the two components of the state process, the nonlinear function $h(x_t)$ is given by

$$h(x_t) = [s^* e^{\gamma x_{1,t} + \beta x_{2,t}} \mathbf{1},\ e^{x_{1,t}} \mathbf{1},\ e^{x_{2,t}} \mathbf{1},\ i^* e^{x_{1,t}} \mathbf{1},\ h^* e^{x_{2,t}} \mathbf{1}]', \tag{34}$$

where $\mathbf{1}$ is an N-dimensional column vector of ones. The nonlinearity of $h(x_t)$ requires using the Unscented Kalman Filter (UKF) to filter out x_t and compute the likelihood function.

C.2 The Unscented Kalman filter

If $h(x_t)$ were linear, the Kalman filter would provide efficient estimates of the conditional mean and variance of the state vector. Let $\hat{x}_{t|t-1} = E_{t-1}[x_t]$ and $\hat{z}_{t|t-1} = E_{t-1}[z_t]$ denote the conditional expectations of x_t and z_t respectively, using information up to $t-1$, and let $P_{t|t-1}$ and $F_{t|t-1}$ denote the corresponding error covariance matrices. Furthermore, let $\hat{x}_t = E_t[x_t]$ denote the expectation of x_t including information at t, and let P_t denote the corresponding error covariance matrix.

The Kalman filter consists of two steps: prediction and update. In the prediction step:

$$\hat{x}_{t|t-1} = \phi_0 + \phi_1 \hat{x}_{t-1}, \tag{35}$$

$$P_{t|t-1} = \phi_1 P_{t-1} \phi_1' + Q, \tag{36}$$

where $\hat{z}_{t|t-1}$ and $F_{t|t-1}$ are in turn given by

$$\hat{z}_{t|t-1} = h_0 + h_1 \hat{x}_{t|t-1}, \tag{37}$$

$$F_{t|t-1} = h_1 P_{t|t-1} h_1' + R. \tag{38}$$

In the update step, the estimate of the state vector is refined based on the difference between observed and predicted quantities:

$$\hat{x}_t = \hat{x}_{t|t-1} + K_t (z_t - \hat{z}_{t|t-1}), \tag{39}$$

$$P_t = P_{t|t-1} - K_t F_{t|t-1} K_t', \tag{40}$$

where K_t is the Kalman gain:

$$K_t = P_{t|t-1} h_1' F_{t|t-1}^{-1}. \tag{41}$$

Since $h(x_t)$ is nonlinear, we adopt the Unscented Kalman Filter (UKF). Rather than approximating the measurement equation linearly (as in the Extended Kalman Filter), the UKF approximates the distribution of the state vector with a deterministically chosen set of "sigma points" that capture the true mean and covariance. These sigma points are then propagated through $h(x_t)$, yielding accurate approximations of the mean and covariance of the data.

Specifically, a set of $2L+1$ sigma points and weights ω^i are selected according to:

$$\chi^0_{t|t-1} = \hat{x}_{t|t-1}, \quad \chi^i_{t|t-1} = \hat{x}_{t|t-1} + \left(\sqrt{(L+\kappa)P_{t|t-1}}\right)_i, \quad \chi^{i+L}_{t|t-1} = \hat{x}_{t|t-1} - \left(\sqrt{(L+\kappa)P_{t|t-1}}\right)_i, \tag{42}$$

with weights

$$\omega^0 = \frac{\kappa}{L+\kappa}, \quad \omega^i = \frac{1}{2(L+\kappa)}, \quad i = 1, \ldots, 2L.$$

The prediction step (equations (37)–(38)) is then replaced by

$$\hat{z}_{t|t-1} = \sum_{i=0}^{2L} \omega^i h(\chi^i_{t|t-1}), \tag{43}$$

$$F_{t|t-1} = \sum_{i=0}^{2L} \omega^i \big(h(\chi^i_{t|t-1}) - \hat{z}_{t|t-1}\big)\big(h(\chi^i_{t|t-1}) - \hat{z}_{t|t-1}\big)' + R, \tag{44}$$

and the update step (equations (39)–(40)) with Kalman gain

$$K_t = \sum_{i=0}^{2L} \omega^i (\chi^i_{t|t-1} - \hat{x}_{t|t-1})(h(\chi^i_{t|t-1}) - \hat{z}_{t|t-1})' F^{-1}_{t|t-1}. \tag{45}$$

Finally, the log-likelihood function is

$$\sum_{t=1}^{T} -\frac{1}{2}\Big[5N\log(2\pi) + \log|F_{t|t-1}| + (z_t - \hat{z}_{t|t-1})' F^{-1}_{t|t-1}(z_t - \hat{z}_{t|t-1})\Big]. \tag{46}$$

C.3 Missing observations handled with UKF

A prominent feature of corporate data are missing observations. In our Compustat panel, 78% of firm-year observations are missing relative to a fully balanced panel. Although the UKF differs from the standard Kalman filter, missing observations can still be handled by the usual method in Kalman filtering.

Suppose there are no missing observations in year t. Then equation (33) holds. If some data in z_t are missing, the key idea is to "select" the observed components using a selection matrix S_t. For example, suppose only the first variable (operating earnings) of the first firm is observed in year t. In that case, $S_t = (1, 0, \ldots, 0)$, which is a $1 \times 5N$ row vector, $M_t = 1$, and $S_t z_t$ is the operating earnings. If all variables of all N firms are available, then S_t is the $5N \times 5N$ identity matrix.

In general, the procedure is: for each t, construct S_t based on the position of observed variables in z_t, pre-multiply both sides of equation (33) by S_t, and run the UKF. Then compute the log-likelihood (46) replacing $5N$ with M_t, the effective number of observations used at time t.

The matrix S_t is time-dependent and allows accommodating missing observations of different variables, as well as firm entry and exit in the panel.

C.4 Mapping W to the data

In our setup, the agent is risk-neutral. In this section, I will consider the general case where the agent has a CARA utility function.

- **Step 1:** CARA Agent optimally consumes γW_t each time.

  The risk-averse agent's continuation payoff under the optimal contract C is:

$$V(C) = \max_{\{\hat{c}_t, e_t \in \{0,1\}\}} E\left[\int_0^\infty e^{-\gamma t} u\left(\hat{c}_t\right) dt\right] (47)$$

$$\text{s.t.} \quad dS_t = rS_t + c_t dt - \hat{c}_t dt, S_0 = 0.$$

  where $u(c) = -e^{bc}/b$, $b \geq 0$ is the agents' risk aversion coefficient, and S_t is their stock of savings. Agents might save by setting the consumption $\hat{c}_t$ below their compensation c_t. A necessary and sufficient condition for no savings is:

$$u'(c_t) = e^{-bc_t} = -b_t \Rightarrow u(c_t) = \gamma V_t \tag{48}$$

- **Step 2:** under risk-neutrality, continuation utility and its certainty equivalent are equal and the agent compensation is $c_t dt = \gamma W_t dt$.
  Define the continuation utility V_t in terms of certainty equivalent: $W_t = -\frac{1}{b\gamma} ln(-b\gamma V_t)$.
  Plug this into the above eq (23), and we get:

$$c_t = \gamma W_t \tag{49}$$

- **Step 3:** the model in the main body is this one in the limiting case where risk aversion goes to zero and we have a smoothing equilibrium.
  Using the martingale representation theorem, the value process of the agents' discounted consumption at time t is:

$$dV_t = \lambda K_t dW_t^x \tag{50}$$

and differentiating equation (22) we get:

$$dV_t + dC_t - \gamma V_t dt = \lambda K_t dW_t^x \tag{51}$$

which is the one used in the main part of the text where $u(c_t) = dC_t$.