

The Coherence Side of Rationality

Theory and Evidence from Firm Plans *

Pamela Giustinelli

University of Padova, Italy, CESifo, and IZA@LISER

Stefano Rossi

Bocconi University, Italy, CEPR, and ECGI

This version: August 2026

Abstract

Using the Duke Survey data on firms' internal plans, we show that over 80% of Chief Financial Officer (CFO) forecasts align closely with simple heuristics taught in MBA textbooks. We introduce forecast coherence—internal consistency across forecasts of jointly determined variables—as a benchmark for evaluating these heuristics. Nearly half of CFOs issue forecasts closely aligned with incoherent heuristics. Such forecasts are associated with predictable reversals in forecast errors, lower firm performance, and underinvestment. We develop a parsimonious model illustrating how reliance on restricted forecasting heuristics can generate incoherence and resource misallocation within firms, consistent with our evidence (*JEL* D84, D22, L2, M2, G32)

*This work was supported by the Baffi Centre on Economics, Finance, and Regulation at Bocconi University. We are grateful to the Editor, Tarun Ramadorai, and the anonymous referees for comments and suggestions that substantially improved the paper. We are also grateful to Alexei Verkhovtsev and Nazrul Shafiq bin Shahril Effendi for outstanding research assistance and to Victor Aguirregabiria, Oriana Bandiera, Nick Barberis, Manudeep Bhuller, Alon Brav, Francesco Corielli, Max Croce, Francesco Decarolis, Alexia Delfino, Aureo De Paula, Xavier Gabaix, Fabian Gouret, John Graham, Marius Guenzel, Max Kellogg, Deborah Kim, Edwin Leuven, Chen Lian, Lance Lochner, Francesca Molinari, Sean Myers, Andras Niedermayer, Marco Ottaviani, Nicola Pavoni, Vincenzo Pezone, Collin Raymond, David Rivers, Matthias Rodemeier, Julien Sauvagnat, Matthew Shapiro, Kelly Shue, Todd Stinebrickner, Johannes Stroebel, Mirko Wiederholt, Wei Xiong, conference participants at the NBER Behavioral Finance Meetings, the Miami Behavioral Finance Conference, the 7th Workshop on Subjective Expectations, the 2nd Northwestern University Former Econometrics Students Conference, the HKU/TLV Finance Forum, the 2023 Western Finance Association Conference, the 2023 International Association of Applied Econometrics Annual Conference, the 2023 Summer Meeting of the European Economic Association, the 2024 CESifo Summer Institute on Expectation Formation, the CSEF-IGIER Symposium on Economics and Institutions, the 5th Joint BoC-ECB-NY FED Conference on Expectations Surveys, the 2025 World Congress of the Econometric Society, and seminar participants at Bocconi University, IWH Halle, Munich TUM, University of Western Ontario, Paris Dauphine University, University of Bristol, University of Oslo, University of Queensland, CY Cergy Paris University, Birkbeck University of London, University of Padova, and Norges Bank for very helpful comments. Send correspondence to Stefano Rossi, stefano.rossi@unibocconi.it.

Coherence—or “the consistency of the elements of the person’s judgment” (Hammond 2007, p. xvi)—is one of the two standards of rationality, together with accuracy.¹ In the context of multidimensional forecasting, coherence requires forecasts of individual variables to incorporate the connections among those variables, whereas accuracy requires individual forecasts not to differ systematically from realizations *ex post*. While forecast accuracy has been widely studied (see, e.g., Tversky and Kahneman [1974] and Benjamin [2019] for a recent review), forecast coherence has received limited attention.

We study forecast coherence in firm plans, internal forecasts over multiple balance-sheet variables, outlining all resource allocation within the firm. When preparing such plans, Chief Financial Officers (CFOs) start from a given output target (“top line”, or sales, forecast), and then proceed to forecast all other items (e.g., input costs).

Internal planning is a challenging multidimensional forecasting problem in a high-stakes setting. Despite its importance, it is still not well understood (Graham 2022). MBA textbooks and case studies offer a series of heuristics, or “rules of thumb” (RoTs), to help CFOs make internal forecasts (e.g., Ruback 2004; Welch 2017; Koller, Goedhart, and Wessels 2020). These rules vary in coherence and sophistication, ranging from forecasting each item on its own to anchoring each forecast to the output target to multivariate rules. To the best of our knowledge, these rules have not previously been assessed. In this paper, we provide a first assessment in terms of forecast coherence and accuracy.

We consider the RoTs as defined in the Welch (2017) taxonomy: (R1) plain growth, (R2) proportion of sales, (R3) economies of scale, (R4) industry based, and (R5) disaggregated. Conceptually, rules (R2)-(R5) reflect some concern for coherence, as they link the forecast of each input’s growth to that of output growth ((R2)-(R4)), and also of other inputs’ growth (R5), whereas (R1) prescribes forecasting each input’s growth by projecting that input’s past growth into the future, implicitly disregarding the input’s relationship with the output and other inputs. Operationally, rules (R1)-(R4) are straightforward, involving a mere sample average (R1) or univariate mean linear regressions ((R2)-(R4)), whereas (R5) is less restricted, using multivariate regression.

¹Coherence can be traced back at least to Aristotle (Fogelin 2003). Tversky and Kahneman (1981, p. 453) write, “there is general agreement that rational choices should satisfy some elementary requirements of consistency and coherence.” See also Sen (1993), Becker (1996), and Posner (2014).

Using balance-sheet and income-statement data from Compustat, we start by documenting that these RoTs give rise to very different forecasts from one another. Next, we assess firm plans by exploiting unique data on simultaneous forecasts of multiple balance-sheet items from the Duke Survey of CFOs of large- and midsize U.S. corporations (Ben-David, Graham, and Harvey 2013; Graham 2022), and linking it to realizations from Compustat.

We find that many CFOs make forecasts that align closely with textbook RoTs and that these rules are not equally good. Some RoTs are associated with more predictable forecast errors, and can be naturally interpreted as more incoherent because they fail to account fully for structural relationships across variables. Firms whose CFO forecasts align more closely with these poorer rules display lower performance and lower investment. These findings are empirical and do not depend on any specific model of belief formation.

We then develop a simple model to organize and interpret our findings. The model provides a benchmark notion of forecast coherence and generates empirical implications for how different RoTs can generate systematically predictable forecast errors and real outcomes, in line with our evidence.

We make three main contributions. First, we document the extent to which (R1)-(R5) describe the forecasts CFOs actually make. We find more than 80% of CFOs provide forecasts closely aligned with (R1)-(R5). The implication is that for a large majority of CFOs it is possible to view their forecasts *as if* they were generated by one of the RoTs, with small discrepancies due to rounding, or small differences in implementation or data.

These findings stop short of demonstrating that CFOs have necessarily used any of the RoTs in Welch (2017). The reason is that our data do not directly elicit the CFOs' forecasting rules, so CFOs could be using some other forecasting process, which turns out to yield forecasts closely aligned with those implied by the RoTs. At the same time, our empirical results suggest that, whether or not CFOs actually use these RoTs, it is useful to think of firm plans in terms of the RoTs in Welch (2017), because those RoTs give rise to forecasts that are significantly different from one another, and because so many CFOs provide forecasts that align closely with those implied by one of the RoTs.

Our second contribution is to document that forecasts in line with different RoTs

covary systematically with firm outcomes in the cross-section. Firms whose CFO forecasts are closest to those implied by (R1) and (R2) have lower firm performance as measured by return on assets (ROA) than those whose forecasts are closest to a multivariate rule. These findings are confirmed using an event-study approach around the dates of CFO appointments. Firms with these same CFO forecasts are also associated with predictable reversals in forecast errors and lower capital expenditures (hence CapEx). Albeit descriptive, this evidence establishes the relevance of the RoTs and suggests the importance of forecast (in)coherence.

Our third contribution is to develop a theory of optimal multidimensional forecasting in firm plans, nesting the five RoTs we study empirically in the first part. The model illustrates formally how different forecasting rules can generate varying degrees of (in)coherence, yielding distinct empirical implications for firm outcomes. In the model, the use of suboptimal forecasting rules provides a novel mechanism generating incoherent plans and, as a result, resource misallocation within the firm.

In our theory, a CFO is given a sales growth target by the CEO and the board of directors. The CFO's task is to minimize total costs while meeting the target. The CFO issues a vector of forecasts over growth rates of firm variables, given knowledge of the firm's technology, which is assumed stable over time and not subject to aggregate shocks. Forecasts are made by minimizing expected inaccuracy under square loss, subject to a coherence constraint provided by the firm's technology. In our benchmark with complete information, we establish that the elements of the optimal forecast vector are coherently linked by the firm's technology parameters, rationalizing a specific version of (R5) as consistent with a first-best benchmark, unlike more restricted RoTs. Our model implies that more restricted RoTs such as (R1) and (R2) are likely to yield lower CapEx forecasts than unrestricted (R5), explaining our evidence of a systematic reversal in forecast errors. In turn, in a world with ex post adjustment costs these lower CapEx forecasts directly translate into underinvestment, consistent with our evidence.

In fact, (R1) is reminiscent of “narrow bracketing” (e.g., [Thaler 1985](#); [Read, Loewenstein, and Rabin 1999](#)), (R5) of “broad bracketing,” and (R2)-(R4) appear to lie in between these two extremes. “Narrow bracketing” refers to the behavior of

decision-makers who, when confronted with multiple related choices (e.g., consumption bundles), make each choice in isolation instead of jointly (“broad bracketing”), thus disregarding how budget constraint and utility function tie choice variables together and obtaining lower utility.

Starting from this observation and drawing on related frameworks of sparse cognition (Gabaix 2014) and narrow thinking (Lian 2021), we model CFOs’ internal plans as an incomplete-information, common-interest game among multiple selves. Each CFO self is in charge of forecasting one balance sheet item, while observing noisy signals about other items. Each CFO self’s imperfect knowledge of the other selves’ forecasts creates intrapersonal frictions in coordinating forecasts of different variables among CFO selves, generating incoherent forecasts. The resultant form of forecast incoherence depends on the signal structure, giving rise to different RoTs.

In our model, “narrow thinking” generates incoherent plans via the use of suboptimal forecasting rules. Incoherent plans, in turn, tend to induce suboptimal resource allocation and, hence, lower profits. In the benchmark case, the multivariate rule (R5) dominates more restricted rules, (R3)-(R4) are intermediate, and (R2)-(R1) perform worse, providing one possible microfoundation for incoherence.

In our data, among the CFOs whose forecasts are closely aligned with (R1)-(R5), 6% are aligned with (R1) and 40% with (R2), with 33% of these exactly equal to (R2). Despite its simplicity and popularity in managerial and consulting textbooks, (R2) is problematic. We show theoretically that under decreasing returns to scale CapEx growth needs to be strictly larger than the sales growth target, implying (R2) is suboptimal. This can also be seen by simply comparing (R2) with (R3) empirically. (R2) can be expressed as a mean regression of each input’s growth on the output’s growth with a zero intercept and a unit slope, which is why we relabel it ‘sales anchoring’. (R3) can be expressed as the same mean regression, but with the intercept and slope actually estimated using Compustat data. The estimated slope is multiplied by the targeted sales growth and added to the estimated intercept. In Compustat, while the slope is close to one, the intercept is positive, large, and significant. Intuitively, the large positive intercept is consistent with replacement and maintenance needs in addition to expansion, implying

CapEx growth typically exceeds sales growth. Importantly, this applies to all industries.

(R1) is also problematic, as it prescribes extrapolating past CapEx growth into the future, whereas in the data CapEx is mean reverting, implying also (R1) yields CapEx growth forecasts that are systematically too low relative to planned sales growth.

Our model also allows to distinguish between forecast coherence and forecast accuracy. Accuracy requires the FEs of each variable to be zero on average, whereas coherence requires the output-input relationship among FEs to mirror that for realizations. This latter theoretical prediction is robust to a large set of technological environments and popular investment frameworks, including Q-theory, time-to-build, investment spikes, financial constraints, irreversibility, and uncertainty. In our data, we reject forecast coherence on average and forecast accuracy for CapEx, but not for sales. Distinguishing coherence from accuracy is important, as under some conditions coherence can be assessed ex ante, and ex ante coherence may aid accuracy ex post.

In our data CFOs are reasonably accurate in forecasting output, reflecting their incentives ([Graham 2022](#)), but much less accurate in forecasting CapEx, likely because their forecast of CapEx growth is incoherently linked to that of sales growth.

Our evidence is also hard to reconcile with several alternative interpretations. Incoherence is distinct from other managerial traits such as optimism and overconfidence (e.g., [Ben-David, Graham, and Harvey \[2013\]](#) and the [Guenzel and Malmendier \[2020\]](#) review): controlling for these traits, we continue to find that incoherence is associated with lower investment, whereas optimism and overconfidence predict the opposite. More generally, our findings are difficult to explain solely by differential managerial control over variables or by firms' optimal responses to idiosyncratic shocks, because CFOs are more accurate about sales than about CapEx, the results are relatively stable across industries, and incoherence is systematically associated with lower firm performance.

Studying coherence is challenging, as it requires referring to “the entire web of beliefs held by the individual” ([Tversky and Kahneman 1974](#), p. 1130) and “to something external to choice behavior (such as objectives, values, or norms)” ([Sen 1993](#), p. 495). Our corporate planning setting is naturally amenable to study forecast coherence, for two reasons. First, executives of large mature firms should have a good sense of their

firm's technology and agree on the firm's objectives. Second, these executives should be optimizing agents, or at least they should be seeking to achieve certain predefined goals.

Incoherence in firm plans provides a possible mechanism for understanding resource allocation within the firm (e.g., Maksimovic and Phillips 2002, 2013; Hoberg and Phillips 2025). Prior literature on managerial beliefs has focused on forecast accuracy of individual variables (e.g., Ben-David, Graham, and Harvey 2013; Graham 2022). Prior work on behavioral bracketing has not studied a firm production setting (e.g., Thaler 1985; Barberis, Huang, and Thaler 2006; Lian 2021). Prior behavioral literature has proposed different mechanisms to generate predictable FEs (e.g., Coibion and Gorodnichenko 2015; Bordalo et al. 2020). Prior behavioral corporate finance literature has focused on other psychological traits, such as optimism and overconfidence (e.g., Baker, Ruback, and Wurgler 2007; Guenzel and Malmendier 2020). We show incoherence differs from both optimism and overconfidence and generates different predictions from the prior literature, which we validate in the data. We discuss the paper's contributions in greater detail in Section 4.

1. Data

We use two main sources of data, one on CFO expectations and one on firm realizations. CFO expectations come from the Duke Survey, run by John Graham and Campbell Harvey. The study surveys 2,000–3,000 CFOs quarterly, asking their views about the U.S. economy and corporate policies, as well as their expectations of future firm performance and operational plans. The typical response rate is 5% to 8%, with most responses arriving within a couple of days. Since the end of the 1990s, the survey has consistently asked respondents their expectations of the future 12-month growth of key firm variables, including revenues, capital expenditures, employment, and earnings.²

Our data comprises 72 quarterly surveys conducted between March 2001 and December 2018. We observe corporate forecasts as a single number per variable, which

²Historical surveys and aggregated responses can be accessed at <https://cfosurvey.fuqua.duke.edu/>.

we interpret as the CFO's expected value, corresponding to the firm's base case scenario.³

Forecasts are elicited for all variables jointly as follows:

Relative to the previous 12 months, what will be your company's PERCENTAGE CHANGE during the next 12 months? (e.g., +3%, -2%, etc.) [Leave blank if not applicable] Revenues: ____; Capital spending: ____; R&D spending: ____; Technology spending: ____; Prices of your product: ____; Earnings: ____; Cash on balance sheet: ____; Number of domestic full-time employees: ____; Wage: ____; Dividends: ____; Advertising: ____; Share repurchases: ____.

Figure A1 of the Internet Appendix shows a screenshot from the survey. It is worth noting that by eliciting all forecasts jointly, the question's format might aid coherence relative to alternative formats eliciting individual forecasts in separate questions.

Firm realizations come from Compustat, which extracts the information from the Security and Exchange Commission (SEC)-required public filing of financial statements. Compustat covers all publicly traded firms across all sectors of the U.S. economy since 1955. We exclude firms with less than \$10 million in total assets or revenue, and we winsorize at the 1% level.

Table A1 of the Internet Appendix shows that, on average, Duke firms are larger in sales and assets, more profitable, and hoard more cash than Compustat firms, but are otherwise similar in terms of market-to-book ratio, investment, and leverage. These patterns concur with prior work using the Duke data (e.g., [Ben-David, Graham, and Harvey 2013](#)).

When matching Duke and Compustat by firm ID, we face four sources of attrition: (1) because of privacy restrictions, not all Duke respondents report their firm ID; (2) not all Duke respondents give forecasts about all variables; (3) some of the variables whose forecasts are elicited in the Duke Survey do not have a precise counterpart in Compustat, namely, technology spending, outsourced employees, health spending, productivity, product prices, and share repurchases; and (4) some of the variables with a counterpart in Compustat do not have full coverage. Chiefly, wages are missing for about 90% of Compustat firms. R&D and advertising expenditures are also missing for many

³For many firms, this is the only scenario leading to fleshed-out forecasts in internal planning. Some firms consider also a downside scenario to plan for contingencies and an upside scenario to lay out stretch goals. However, the latter do not usually lead to fleshed-out forecasts ([Graham 2022](#), p. 1997).

Compustat firms.⁴

Table A2 of the Internet Appendix reports summary statistics on 12-month-ahead growth forecasts (panel A), followed by the corresponding realizations (panel B), in the matched Duke-Compustat sample.⁵ Comparing the two panels shows that while CFOs are on average slightly optimistic about sales (output), the empirical medians of forecasts and realizations are quite close to each other, consistent with the [Graham \(2022\)](#) observation that CFOs care about getting sales forecasts right. Conversely, CFOs' forecasts of capital expenditures are systematically too low, with the distribution of forecasts shifted to the left relative to the distribution of realizations.

2. Empirical Analysis

When preparing firm plans, CFOs typically start by issuing forecasts of sales growth for a number of years, reflecting targets set by the CEO and board members, and then proceed to forecast all other balance-sheet variables ([Graham 2022](#)). To help CFOs with this multidimensional forecasting problem, managerial textbooks and case studies propose simple methods, surveyed in the following taxonomy by [Welch \(2017, pp. 593–4\)](#):

- (R1) A **plain growth** forecast, separately projecting past growth rates of each item; henceforth 'narrow bracketing' rule. [Welch \(2017\)](#) implements this rule by computing for each item the average of the two most recent annual growth rates.
- (R2) A pure **proportion of sales** forecast, forecasting each item as a fixed proportion of sales; henceforth 'sales anchoring' rule. [Welch \(2017\)](#) implements this rule by assigning to each item the same growth rate as that forecasted for sales. Thus, this rule can be expressed as a mean regression of each item's growth on sales' growth with a zero intercept and a unit slope.

⁴Points 1 and 2 raise the possibility of nonrandom selection. However, under the assumption that missing forecasts reflect lack of knowledge, our respondents are likely positively selected from those more likely to provide complete and high-quality answers. Points 3 and 4 imply that our analysis of FEs needs to be limited to items with full coverage in both Duke and Compustat. Finally, the matched Duke-Compustat sample mostly refers to the first part of the sample period, until 2011Q4. This is not a problem for us, as we conduct most of our regression analysis in the prefinancial crisis period.

⁵For some firm-year pairs there might be multiple CFO forecasts, as we match Duke and Compustat by firm ID. Table A3 of the Internet Appendix reports the same statistics in the full Compustat population.

- (R3) An **economies-of-scale** forecast, positing for each item a fixed component and a variable component, the latter itself a proportion of sales' growth. [Welch \(2017\)](#) implements this rule by estimating a best linear predictor under square loss of each item's growth on contemporaneous sales growth using Compustat data. The estimated regression intercept is the fixed component and the estimated slope multiplied by the forecasted sales growth is the variable component.
- (R4) An **industry-based** forecast, drawing on information from other firms in the same industry. [Welch \(2017\)](#) implements this rule exactly as (R3), but using only data on firms in the same industry as the firm under consideration.
- (R5) A **disaggregated** forecast, allowing each item to comove not only with sales but also with other items. [Welch \(2017\)](#) implements this rule by using a 'kitchen sink' specification with all Compustat variables.

There is no consensus on which of the above rules, if any, constitutes best practice. In their best-selling guide to valuation, [Koller, Goedhart, and Wessels \(2020\)](#) advocate for (R2), writing “net Property, Plant and Equipment should be forecast as a percentage of revenues” (p. 286).⁶ Similarly, [Ruback \(2004\)](#) advocates for (R2), stating “capex should be proportional to sales” and Harvard Business School case studies provide solutions in which CapEx growth forecasts equal sales growth forecasts, in line with (R2), and also suggest (R1) and (R4) (e.g., [Luehrman and Heilprin 2009](#); [Stafford and Heilprin 2011](#)). Proponents of other rules include [Titman and Martin \(2016\)](#) for (R3) and [Holthausen and Zmijewski \(2020\)](#) for (R5). This literature lacks a framework to evaluate these methods and establish whether they are equivalent to one another or, if not, which ones perform best.

This section provides such a framework. Subsection 2.1 examines the extent to which (R1)-(R5) describe the forecasts CFOs make. Subsection 2.2 proposes a measure of (in)coherence. Subsection 2.3 presents tests of coherence and accuracy based on forecast errors. Subsection 2.4 examines whether forecasts aligned with RoTs are associated with

⁶In turn, [Koller, Goedhart, and Wessels \(2020\)](#) indicate that depreciation should be forecasted as a proportion of net Property, Plant, and Equipment (PPE), and that capital expenditures should be calculated by summing the projected increase in net PPE to depreciation (pp. 286-7). As a result of this chain of calculations, capital expenditures is forecasted as a percent of revenues, and thus the forecast of capital expenditures growth should equal the forecast of sales growth, as in (R2).

predictable forecast errors. Subsections 2.5 and 2.6 examine the associations between incoherence (RoTs) and performance and CapEx. Subsection 2.7 presents an event study of changes in incoherence (RoTs) and firm outcomes around CFO appointments.

2.1 *RoTs and their distance from CFO forecasts*

We begin in Table 1 with a dual purpose. First, Table 1 describes how to implement specific RoTs following Welch (2017) and using Compustat data; second, Table 1 provides a first assessment of the RoTs, exclusively based on Compustat data.

The first row of Table 1 reports the results of a regression designed to assess (R1), which dictates to forecast future CapEx growth as the average of the past two CapEx growth rates. We estimate, $\text{CapEx Growth}_{i,t \rightarrow t+1} = \alpha + \beta \text{Avg CapEx Growth}_{i,t-2 \rightarrow t} + \varepsilon_{i,t+1}$, where $\text{Avg CapEx Growth}_{i,t-2 \rightarrow t}$ denotes the average CapEx annual growth rate from $t - 2$ to t . If $\alpha = 0$ and $\beta = 1$, (R1) would yield on average an accurate forecast of CapEx growth. In the data, however, $\hat{\alpha} = 0.339$ and $\hat{\beta} = -0.101$, precisely estimated and statistically different from 0 and 1, respectively, implying CapEx growth is mean-reverting, which is not captured by (R1).

The second row of Table 1 reports the results of a regression designed to implement (R3) and assess (R2). We estimate, $\text{CapEx Growth}_{i,t \rightarrow t+1} = \alpha + \beta \cdot \text{Sales Growth}_{i,t \rightarrow t+1} + \varepsilon_{i,t+1}$. We find $\hat{\alpha} = 0.230$ and $\hat{\beta} = 1.080$, precisely estimated. Thus, a firm wishing to achieve a 5% sales growth rate would, under (R3), forecast a CapEx growth rate of $1.080 \times 5\% + 23\% \approx 28\%$. These estimates imply that (R2), which assumes $\alpha = 0$ and $\beta = 1$, provides much lower CapEx growth forecasts than (R3). Indeed, the same firm wishing to achieve a 5% sales growth rate would forecast under (R2) a CapEx growth rate of only 5%. The discrepancy between the two rules is driven by the intercept: equal to zero according to (R2), positive and large according to (R3). The estimated slope of 1.080 under (R3) is statistically indistinguishable from 1 assumed by (R2). Relative to (R3), (R2) underpredicts the growth rates of other inputs as well, though to a smaller extent. For example, a firm aiming at a 5% sales growth should under (R3) forecast a 11% advertising growth and a 8% wage growth.

The subsequent 10 rows implement (R4) by estimating the same regression separately

for each of 10 industries, defined by one-digit SIC codes.⁷ The intercept ranges from 0.163 to 0.369, and the slope ranges from 0.881 to 2.037, and both are statistically significant in all 10 industries. These results confirm that in all industries (R2) implies a lower CapEx growth forecast relative to sales growth than in the data.

The last row reports results from a multivariate regression implementing (R5), $\text{CapEx Growth}_{i,t \rightarrow t+1} = \alpha + \beta \cdot \text{Sales Growth}_{i,t \rightarrow t+1} + \theta \cdot \text{Earnings Growth}_{i,t \rightarrow t+1} + \varepsilon_{i,t+1}$. Earnings forecasts are widely reported and earnings realizations are widely available, unlike other variables. We find $\hat{\alpha} = 0.229$, $\hat{\beta} = 1.067$, and $\hat{\theta} = 0.017$, precisely estimated.

As a next step, we measure the distance between the CFO forecast reported in the Duke data and the forecasts implied by each RoT. First, we compute five values, each equal to the orthogonal distance between the CFO's forecast and that implied by each RoT. Second, we compute the minimum among these five values. For each CFO, i , we refer to the distance-minimizing RoT, $\tau^i \in \{1, \dots, 5\}$, as to the CFO's 'type'.⁸

Panel A of Table 2 shows that, among the 396 CFOs for whom we observe their identity and the joint forecasts and realizations of all variables, a plurality of about 40% makes a forecast that is closest to that implied by (R2), 27% of whom appear to use exactly (R2). Another 7.6% of CFOs are closest to (R1). Thus, about 47% of CFOs make forecasts closest to an incoherent RoT. We also find that about 10% of CFOs make forecasts closest to (R3), 28% to (R4), and 15% to (R5). These results underscore the large heterogeneity in CFO forecasts, possibly reflecting the challenges of making coherent forecasts and the absence of a consensus on which RoTs to use.

The average minimum distance of 0.033 is reasonably small, while masking substantial heterogeneity as revealed by a standard deviation of 0.059. Small discrepancies between actual and minimum-distance forecasts may reflect rounding or truncation in CFOs' reports and/or small differences in rule implementation between CFOs and us, for example, the use of a different number of past observations for (R1) or a different sample

⁷Agriculture, forestry, fishing; Mining; Construction; Manufacturing; Transportation, communication, utilities; Wholesale trade; Retail trade; Finance, insurance, real estate; Services; Public administration.

⁸To focus on coherence and minimize confounding with accuracy, we compute the distance relative to the sales forecast, as CFOs are more accurate at forecasting sales than any other item, as we will show later. In the Internet Appendix, we present multiple robustness checks. Table A4 implements (R5) using advertising instead of earnings. Tables A5 and A6 report the distributions of RoTs, when the distance is measured relative to earnings and CapEx, respectively.

size for (R3)-(R5). Larger discrepancies raise the possibility forecasts reflect different methods. To investigate this possibility, we compute the fraction of CFOs for whom their actual forecast falls outside of a prespecified interval around the minimum-distance one. We find that the actual forecast falls within a ± 0.05 interval around the forecast implied by one of the five RoTs for 81.8% of CFOs and within a ± 0.025 interval for close to 60% of CFOs. By way of comparison, the standard deviation of sales growth forecasts is 0.271 (see panel A of Table A2). Hence, the widths of the ± 0.05 and ± 0.025 intervals are about one-third and one-sixth, respectively, of that standard deviation. These figures suggest our classification is plausible.

To account for the possibility that some CFOs use a different method, in panel B we report an alternative categorization in which we assign to a residual category, (R6), the 18.2% CFOs whose actual forecasts are outside of a ± 0.05 interval around the forecasts implied by (R1)-(R5). Of the CFOs still categorized as (R1)-(R5), 6.2% make forecasts closely aligned with (R1), 39.8% with (R2), 11.4% with (R3), 27.5% with (R4), and 15.1% with R5. We use this alternative categorization in robustness analyses.

2.2 Incoherence

To construct a measure of incoherence, we require a reference forecasting rule that captures the contemporaneous linear relationship among multiple balance-sheet items. Among the rules in Welch (2017), (R5) provides the least restrictive linear specification, allowing each item to comove with sales and at least one additional item. The other rules can be viewed as restricted versions of this specification: (R3)-(R4) omit additional inputs or focus on subsamples, (R2) imposes a zero intercept and unit slope, and (R1) projects each item in isolation. Because (R5) spans the space of linear joint forecasts considered in this taxonomy, we use its implied hyperplane as a reference and define incoherence as the orthogonal distance of a CFO's forecast from that multivariate benchmark.⁹

We implement (R5) as, $\Delta \log I_t = \alpha + \beta \cdot \Delta \log Y_t + \theta \cdot \Delta \log \Psi_t + \varepsilon_t$, where Ψ denotes

⁹By doing this, we do not assume that (R5) is the optimal rule. Rather, we use it as a statistical reference because it represents the most general linear relationship among the items considered in the Welch (2017) taxonomy. Other rules can be expressed as restrictions of this multivariate form. Our incoherence measure captures the extent to which a CFO's forecasts depart from this benchmark.

profits. This equation corresponds to the standard log-linear investment relationship derived from cost minimization with quadratic adjustment costs, as we show in the Internet Appendix A.1 (e.g., Hayashi 1982; Abel and Eberly 1994). It represents the contemporaneous relationship among CapEx, output, and earnings growth. We measure output growth, $\Delta \log Y_t$, with sales growth, as is standard in the literature (e.g., Syverson 2011). This allows us to bypass the paucity of data on labor costs available in Compustat.

To focus on coherence while minimizing confounding with forecast accuracy, we measure distance in the sales-forecast dimension. We therefore rewrite the same hyperplane in terms of sales growth as, $\Delta \log Y_t = \alpha^Y + \beta^Y \cdot \Delta \log I_t + \theta^Y \cdot \Delta \log \Psi_t + \varepsilon_t^Y$, and define incoherence as the orthogonal distance between the CFO's forecast of sales growth and the sales-growth forecast implied by this hyperplane:

$$\text{Incoherence}_{i,t} = \frac{|\Delta \mathbb{F}_{i,t}[Y_{i,t+1}] - \hat{\beta}^Y \Delta \mathbb{F}_{i,t}[I_{i,t+1}] - \hat{\theta}^Y \Delta \mathbb{F}_{i,t}[\Psi_{i,t+1}] - \hat{\alpha}^Y|}{\sqrt{1 + (\hat{\beta}^Y)^2 + (\hat{\theta}^Y)^2}}, \quad (1)$$

where $\Delta \mathbb{F}_{i,t}[x_{i,t+1}] = \log \mathbb{F}_{i,t}[x_{i,t+1}] - \log \bar{x}_{i,t}$ denotes the forecast of the growth rate of a generic variable x , and $\hat{\alpha}^Y$, $\hat{\beta}^Y$, and $\hat{\theta}^Y$ are the estimated coefficients from the sales-side representation of the multivariate hyperplane implied by (R5). The denominator normalizes the numerator to produce the shortest orthogonal distance from the forecast point to that hyperplane, as we show in the Internet Appendix A.2.

We have 396 CFOs in our main sample. When using advertising instead of earnings growth (as in column 8 of Table A7 of the Internet Appendix), we obtain similar results, although based on 177 CFOs only. Similarly, using wages or R&D instead of earnings growth, or more than two regressors, results in broadly similar patterns, though with even smaller samples of 52 or fewer CFOs. Rather than the specific regressors, what turns out to be crucial is moving from a univariate to a multivariate specification, that is, that forecasters account for the contemporaneous relationship among multiple items.

Because different forecasting rules impose different restrictions on the joint behavior of balance-sheet items, one would expect incoherence to vary systematically across rules. To examine this variation, we regress our incoherence measure in (1) on binary indicators for (R1)-(R4), with (R5) as the reference category.

Table 3 presents the results and shows that (R1) has the largest incoherence relative to (R5), followed by (R2). The average differences in incoherence of (R3) and (R4) relative to (R5) are much smaller in magnitude.

Robustness analyses yield similar results. When using advertising instead of earnings to implement (R5), we still find (R1) is most distant from (R5), followed by (R2), and both are strongly statistically significant (see Table A8). When using the alternative categories (R1)-(R6) in Table A9, we find as expected that (R6) is the most distant from (R5) and, more remarkably, that the relative ranking of (R1)-(R4) is unchanged, with only (R1) and (R2) statistically significantly different from (R5).

Table 4 reports sample characteristics and investigates heterogeneity in ex ante incoherence by CFO and firm characteristics. In panel A, 45% of CFOs have an MBA and 9% are women. On average, CFOs are 50.4 years old and have been on the job for 4.3 years. Average firm size is US\$2.5 billion of sales and 64% of firms pay a dividend. These figures are in line with those in [Ben-David, Graham, and Harvey \(2013\)](#).

Panel B shows incoherence is lowest at the intermediate age range, 41-50, suggesting experience may help CFOs form more coherent forecasts, and coherence might decline at older ages. Incoherence is unrelated to optimism and miscalibration, consistent with the idea from psychology that incoherence and overconfidence are different constructs. Having an MBA does not correlate with incoherence, likely reflecting that, on the one hand, some RoTs are quite simple and CFOs may come up with them on their own and, on the other hand, there is no consensus in MBA textbooks on which rule is best; although in practice, as shown in Table 3, different rules perform differently in terms of ex ante coherence. Finally, CFOs of larger firms and firms paying dividends display lower incoherence, although the statistical significance is marginal. These results suggest that in stable or predictable environments it might be easier to make (more) coherent forecasts.¹⁰

¹⁰CFO tenure does not correlate with incoherence. This, combined with the fact that longer tenured CFOs delegate less ([Graham, Harvey, and Puri 2015](#)), suggests that incoherence is unrelated to delegation of authority.

2.3 Tests of coherence (and accuracy) using forecast errors

Coherence imposes intuitive restrictions on how forecast errors of output and inputs must relate to one another. Consider the following empirical formulation of a generic production function,

$$y_{t+1}^f = \alpha + \alpha^f + \sum_{i=1}^n \beta_i^f x_{i,t+1}^f + \sum_{i=1}^n \sum_{s=0}^t \delta_{i,t-s}^f x_{i,t-s}^f + \sum_{j=1}^m \sum_{s=0}^t \gamma_{j,t-s}^f z_{j,t-s}^f + \varepsilon_{t+1}^f, \quad (2)$$

where f indexes firms, i inputs, t time, s lags, and j any relevant state variable (e.g., inventory, cash). The loadings of inputs and state variables, β_i^f , $\delta_{i,t-s}^f$, and $\gamma_{j,t-s}^f$, may vary across firms and over time, allowing for a wide array of lead-lag relationships. For example, firms in short-term industries will load more on inputs and state variables in recent years, whereas firms in long-term industries will load more on inputs and state variables in more distant years. The variables could be in levels, growth rates, logarithms, or any other transformation that preserves linearity in parameters. Note this formulation allows for many empirically relevant cases, including standard Q-theory, investment spikes, time-to-build considerations, financing costs and other state variables, and allowing for a quite general lead-lag relationship between investment and output.

Coherent forecasts of firm-specific variables should be cross-sectionally related as are their realizations in Equation (2), that is,

$$\mathbb{F}_t [y_{t+1}^f] = \alpha + \alpha^f + \sum_{i=1}^n \beta_i^f \mathbb{F}_t [x_{i,t+1}^f] + \sum_{i=1}^n \sum_{s=0}^t \delta_{i,t-s}^f x_{i,t-s}^f + \sum_{j=1}^m \sum_{s=0}^t \gamma_{j,t-s}^f z_{j,t-s}^f. \quad (3)$$

Subtracting (3) from (2) *within firms* yields

$$\mathbb{F}\mathbb{E}_t [y_{t+1}^f] \equiv y_{t+1}^f - \mathbb{F}_t [y_{t+1}^f] = \sum_{i=1}^n \beta_i^f \mathbb{F}\mathbb{E}_t [x_{i,t+1}^f] + \varepsilon_{t+1}^f, \quad (4)$$

Differently from Equation (3), Equation (4) does not depend on any firm-level heterogeneity known to, or predictable by, the CFO at the time of forecast—but likely unobserved by the econometrician—as such heterogeneity gets differenced away when subtracting the forecast from the realization.

This discussion suggests two intuitive restrictions on FEs of contemporaneous inputs

and output. First, as long as inputs contribute positively to output, that is, input loadings are positive, one would expect FEs of output and each input to be positively associated. Second, as long as no individual input has increasing returns, that is, input loadings are not larger than one, one would expect forecast errors to lie below the 45-degree line. Together, these restrictions imply the shaded areas in Figure 1.

Figure 2 plots the joint distribution of FEs of Sales and CapEx. While the regression slope is positive as expected, 42% of CFOs have FEs of Sales and CapEx of opposite sign, violating the first restriction of positive loadings. An additional 10.4% of FE pairs imply an input loading larger than one, thus violating the second restriction. Therefore, more than 50% of CFOs report contemporaneous FEs that violate intuitive coherence bounds. We find similar results in all 10 industries, consistent with the idea that examining contemporaneous FEs is informative about coherence regardless of the lead-lag relationships or the short- versus long-term outlooks of different industries. We also find similar results for additional FE pairs, including sales and earnings.

Next, we assess Equation (4) by regressing FEs of output on FEs of inputs,

$$\mathbb{F}E_t \log y_{t+1}^f = \alpha + \beta_1 \cdot \mathbb{F}E_t \log x_{1,t+1}^f + \beta_2 \cdot \mathbb{F}E_t \log x_{2,t+1}^f + \varepsilon_{t+1}^f, \quad (5)$$

whereby under the null of forecast coherence one would expect the constant to equal zero and the slope of the FE of each input to equal the loading of the corresponding input in (2). In the Cobb-Douglas benchmark with capital stock and labor as inputs, this implies $\alpha = 0$, $\beta_1 = a$, and $\beta_2 = b$, where a and b are the capital and labor shares,¹¹ which we compute using data from the Bureau of Economic Analysis.¹²

This framework further allows to distinguish between forecast accuracy and coherence, as accuracy can be assessed by testing for each variable that the mean of the

¹¹The interpretation $\beta_1 = a$ holds exactly when the capital-side input is capital stock. In our empirical implementation, we use capital expenditures, so the mapping from theory to Equation (5) is reduced-form: under $K_{t+1} = K_t(1 - \delta) + I_t$, forecast errors in CapEx proxy locally for forecast errors in the capital-side input, but the exact coefficient need not equal the capital share.

¹²To implement the tests, we need $\mathbb{F}E_t \log x_{1,t+1}$ and $\mathbb{F}E_t \log y_{t+1}$, which are not directly observed in our data as the forecasts were not elicited in logs. To address this issue, we use $\mathbb{E}_t \log x_{t+1} \approx \log \mathbb{E}_t x_{t+1} - \frac{1}{2} V_t \log x_{t+1}$, a transformation written for a generic input x , and $V_t \log x_{t+1}$ is the conditional variance of $\log x$. Hence, $V_t \log x_{i,t+1} = \sigma_i^2 (1 - \gamma_i^2)$ for $i = 1, 2$, where σ_i^2 is the unconditional variance and γ_i the coefficient of an AR(1) regression of $\log x_{i,t}$. The conditional variance of y is $V_t \log y_{t+1} = a^2 \sigma_1^2 + b^2 \sigma_2^2$.

FEs is zero, $\text{Avg } \mathbb{E}_t [y_{t+1}^f] = 0$ and $\text{Avg } \mathbb{E}_t [x_{i,t+1}^f] = 0$ for each input i . Equivalently, accuracy can be tested by regressing realizations on forecasts for each variable.¹³

Table 5 reports the results. In columns 1 and 2, we test for accuracy of CapEx and sales growth. In column 1, the constant is statistically different from zero at 10%; in column 2 the constant is not statistically different from zero. Hence, CFOs are on average inaccurate about CapEx and accurate about sales. The latter is consistent with the [Graham \(2022\)](#) observation that top executives care the most about sales forecasts.

In column 3 we test for coherence. We find $\hat{\alpha} = 0.046$, statistically different from 0 at 5%; $\hat{\beta}_1 = 0.113$, marginally significant at 10%; and $\hat{\beta}_2 = 0.023$, not statistically different from zero. Hence, forecast errors of output and inputs do not satisfy the joint restrictions implied by coherence, and we reject the null of coherence.

The coherence test in column 3 requires FEs on wages, implying we could perform it on 51 observations only, raising issues of statistical power. In column 4, we implement the regression test without wages. We now find $\hat{\alpha} = -0.004$, not statistically different from zero, and $\hat{\beta}_1 = 0.135$, precisely estimated. This reduced-form specification yields a similar qualitative pattern, although omitting wages weakens the direct structural interpretation of the coefficient. In column 5 we estimate (4) without CapEx and find the coefficient on wages to remain statistically insignificant.

In sum, CFO forecasts are on average accurate for sales, inaccurate for CapEx, and incoherent across the two. These results indicate that accuracy and coherence are not only distinct concepts in theory, as illustrated in Figure 3, but also in our data. These conclusions are confirmed in individual-level tests based on stronger assumptions, which we derive, implement, and report in Internet Appendix C.

¹³Under rationality and no aggregate shocks, one expects a unit slope and a zero intercept. There is a long tradition of rational expectations tests of this type (e.g., [Muth 1961](#); [Lovell 1986](#); [Pesaran and Weale 2006](#); [D'Haultfoeuille, Gaillac, and Maurel 2021](#); [Crossley et al. 2024](#); [Giustinelli and Shapiro 2024](#)).

2.4 *Incoherence, rules of thumb, and predictable forecast errors*

We next examine whether incoherent forecasts are associated with systematically predictable forecast errors. First, we estimate the following specification,

$$\mathbb{E}_t \text{CapEx Growth}_{t+1}^f = \beta_0 + \beta_1 \cdot \text{CapEx Growth}_t^f + \beta_2 \cdot \text{Sales Growth}_t^f + \beta_3 \cdot \text{Earnings Growth}_t^f + \varepsilon_{t+1}^f.$$

Table 6 reports the results. In columns 1 and 2, we find $\hat{\beta}_1 < 0$, statistically significant, indicating high current CapEx growth predicts low subsequent FEs. Hence, CFOs overreact to realizations of CapEx growth when forecasting future CapEx growth.

Next, we examine whether this pattern differs by RoT. In columns 3 and 4 we find the largest overreaction for CFOs whose forecasts are closest to (R1), (R2), and (R4), with (R2) statistically significant at 5% level.

These results differentiate our paper from recent literature that links specific deviations from perfect rationality to systematic forecast errors. [Coibion and Gorodnichenko \(2015\)](#) show how investor inattention generates underreaction. [Bordalo et al. \(2020\)](#) show how diagnostic expectations generate overreaction. We also document overreaction, but we link it to a different mechanism, namely, incoherent forecasting heuristics, as in Table 6. We will discuss later in Section 3 one possible theoretical microfoundation of this result.

So far, we have distinguished forecast coherence from accuracy, and shown that incoherence generates predictable FEs. Next, we test whether incoherence and incoherent heuristics are associated with real effects on performance and investment.

2.5 *Incoherence, rules of thumb, and firm performance*

We now examine whether corporate performance varies with managerial incoherence, by estimating the regression,

$$\text{ROA}_{i,j,t} = \alpha + \lambda_j + \delta_t + \beta \cdot \text{Incoherence}_{i,j,t} + \theta \cdot X_{i,j,t} + \varepsilon_{i,j,t},$$

where i indexes CFO-firm pairs, j indexes industries, t indexes time, the dependent variable $\text{ROA}_{i,j,t}$ is the percent return on the firm's assets, λ_j are industry fixed effects, δ_t

are survey fixed effects, and $X_{i,j}$ includes firm-level regressors (firm size, market-to-book, and dividends) and CFO-level regressors (miscalibration and optimism, measured at both short- and long-term horizons). If deviations from coherence in internal plans reflect misallocation of inputs, higher incoherence should be associated with lower firm performance. Accordingly, we expect $\beta < 0$.

We conduct the estimation over 2001-2007 to avoid the effect of the financial crisis, an aggregate shock which affected managerial beliefs (Boutros et al. 2025), and to focus on the peak of the “great moderation”, when aggregate volatility was not a concern. We compute bootstrap standard errors and cluster them at the firm level. Given the lack of exogenous variation in incoherence, the results should be interpreted as correlational.

Table 7 shows the results. In column 1, a one-standard-deviation increase in incoherence (0.079 from Table 3) is associated with a 3 percentage points (p.p.) lower ROA, significant at the 5% level. Results are stable when we further condition on CFO miscalibration and optimism (columns 2 and 3), industry and survey fixed effects (column 4), and firm-level regressors (columns 5-8). Thus, while sample size changes across columns reflecting availability of regressors, the estimated negative association between CFO incoherence and firm performance remains quite stable in magnitude and statistically significant.

Next, we assess the extent to which this result reflects different RoTs by regressing firm performance on RoT indicators and other regressors. Consistent with the ranking documented in Table 3, in Table 8 we find that firms whose CFO forecasts are closest to (R1) have a 5- to 6-p.p. lower ROA relative to (R5). These are large differences in economic terms. We also find firms whose CFO forecasts are closest to (R2) or (R3) have a 2.5- to 4.5-p.p. lower ROA relative to (R5), whereas (R4) are indistinguishable from (R5).

Table A10 in the Internet Appendix reports similar results when using an alternative definition of (R5) based on advertising growth.

In a robustness analysis using the (R1)-(R6) categorization, we find that firms whose CFOs are classified as (R6) have a 4-p.p. lower average ROA and results for (R1)-(R4) firms are very similar to those presented in Table 8. Therefore, CFOs whose forecasts

are quite distant from those implied by all the five rules we study are unlikely to use a more sophisticated or “better” approach, further supporting our focus on the [Welch \(2017\)](#) taxonomy and our version of (R5) as a benchmark rule. In sum, firm performance is negatively associated with incoherence, and is on average lowest for firms whose CFO forecasts are closest to (R1) or (R2).

2.6 *Incoherence, rules of thumb, and corporate investment*

If downward-biased capital expenditure forecasts are only partially corrected ex post when firms implement their plans, firms relying on incoherent rules should exhibit lower realized investment. In Table 9, we assess this prediction by estimating,

$$Y_{i,j,t} = \alpha + \lambda_j + \delta_t + \beta \cdot \text{RoT}_{i,j,t} + \theta \cdot X_{i,j,t} + \varepsilon_{i,j,t},$$

where the indexes i , j , and t are defined as in the earlier regression for ROA; $\text{RoT}_{i,j,t}$ is a vector of indicators for (R1)-(R4); $Y_{i,j,t}$ is CapEx over assets (columns 1-3), and $X_{i,j,t}$ is a vector of covariates including firm size, market-to-book ratio, and dividends.

In columns 1 and 2, (R1) and (R2) are associated with 1.3- to 1.6-p.p. lower CapEx relative to (R5). The difference is large in economic terms and for (R2) statistically significant. These results continue to hold after conditioning on miscalibration and optimism. We confirm the [Ben-David, Graham, and Harvey \(2013\)](#) finding that miscalibration and optimism are positively (rather than negatively) associated with investment spending, underscoring that miscalibration, optimism, and incoherence are distinct constructs. In sum, (R1) and (R2) are associated with lower investment relative to (R5).

2.7 *Change in firm behavior when CFOs take office*

The negative association between performance and incoherence may reflect different possibilities. Higher incoherence may lead to lower investment and performance. Or, lower investment may induce CFOs to make incoherent forecasts, for example, to forecast too high a sales growth. Finally, more incoherent CFOs might be selected by—or self-select

into—firms with lower investment spending and worse performance.

To make a step toward causality and shed light on its direction, we study how firm performance and investment evolve in the years around a CFO’s hiring. We extract the dates when CFOs join firms from Execucomp and Boardex data and supplement this information by hand-collecting data from 10-K filings. A CFO is considered to take office in a firm when he or she first signs the financial reports. We match performance, investment, and characteristics from Compustat for the year of taking office. We use as dependent variables the within-firm differences in ROA and investment between the year before and after the CFO’s taking office.

Table 10 presents the results. Performance declines after the appointment of incoherent CFOs, especially those with forecasts closest to (R1). (R1) is also associated with a 2.2-p.p. lower investment intensity. While not unambiguously causal, these results align with the evidence documented above.

3. A Simple Model of Forecast Coherence

We present a model to illustrate forecast coherence and interpret our empirical results. Subsection 3.1 introduces the benchmark model and shows incoherent RoTs yield lower performance and overreaction. Subsection 3.2 shows incoherent RoTs yield underinvestment. Subsection 3.3 discusses one possible microfoundation of incoherence.

3.1 *A benchmark model of optimal corporate forecasts*

Consider a firm producing output y_t at time t using n inputs, where x_i is the quantity of input i , p_i is its price, and p_y the output price, which we normalize to one. To map the theory into our data, in the main text we develop the Cobb-Douglas case with two inputs (K, L) , $y = (x_1^a \cdot x_2^b)$, and we consider the more general CES case in Internet Appendix A.6.

The CFO’s task at any time t is to choose inputs to minimize costs subject to meeting a given sales growth target and taking input prices as given. While capital stock follows, $K_{t+1} = K_t(1 - \delta) + I_t$, where δ is the depreciation rate and I_t is capital expenditures, at

any time t at which the CFO forecasts 1-year-ahead variables, the current stock of capital, K_t , is given, and all that matters for expected 1-year-ahead profits are the expected output and the capital (and labor) expenditures. As a result, the CFO aims to produce input growth forecasts that are both individually accurate and jointly coherent, in the sense that they are consistent with the firm's production technology.

We assume the technology is stable over the forecasting horizon and abstract from aggregate shocks, so factor-augmenting productivities are constant and normalized to one. We further maintain standard producer-theory assumptions: firms take input and output prices as given and choose production plans to minimize costs.

The CFO is in charge of internal plans, which consist of forecasts of growth rates of multiple balance-sheet variables. We denote the forecast of sales growth at time t by $\Delta \mathbb{F}y_{t+1} = \log \mathbb{F}y_{t+1} - \log \bar{y}_t$, and the forecasts of the inputs' growth rates at time t by $\{\Delta \mathbb{F}x_{i,t+1} = \log \mathbb{F}x_{i,t+1} - \log \bar{x}_{i,t}\}_{i=1,2}$, where $\bar{\eta}_t$ denotes the realization of the generic variable η . It is important to note that the forecasters observe realizations of own and other firms' output and inputs, for example, by accessing Compustat. On the other hand, the forecasters observe only their own past and current forecasts, but not those of other firms.

We consider the problem of a CFO who takes $\Delta \mathbb{F}y_{t+1}$, set by the CEO and board of directors, as given, and forecasts the growth rate of input expenditures under square loss,

$$\min_{\Delta \mathbb{F}x_{1,t+1}, \Delta \mathbb{F}x_{2,t+1}} \mathbb{E}[(\Delta x_{1,t+1}^* - \Delta \mathbb{F}x_{1,t+1})^2 + (\Delta x_{2,t+1}^* - \Delta \mathbb{F}x_{2,t+1})^2 \mid \Omega_t],$$

where Ω_t denotes the CFO's information set at t , which includes knowledge of the production function, and $\{\Delta x_{1,t+1}^*, \Delta x_{2,t+1}^*\}$ denote the cost-minimizing inputs growth needed to obtain the sales growth target, $\Delta \mathbb{F}y_{t+1}$.

In words, the CFO chooses the vector of input-growth forecasts that minimizes expected inaccuracy, subject to a coherence constraint that forecasts be consistent with the firm's technology and sales growth target. This formulation maps the CFO's forecasting problem into a standard cost-minimization problem. An analogous analysis would obtain under profit maximization. Square loss lets the model nest RoTs (R1)-(R5).

3.1.1 Optimal forecasts.

We begin by establishing two results about coherent forecasts and forecast errors.

Lemma 1 (Coherent forecasts under Cobb-Douglas). *Forecast coherence requires that forecasts of output and inputs, $\mathbb{F}_t[y_{t+1}]$, $\mathbb{F}_t[x_{1,t+1}]$, and $\mathbb{F}_t[x_{2,t+1}]$, satisfy,*

$$\begin{aligned}\mathbb{F}_t \log[y_{t+1}] &= a \cdot \mathbb{F}_t \log[x_{1,t+1}] + b \cdot \mathbb{F}_t \log[x_{2,t+1}], \\ \Delta \mathbb{F}_t[y_{t+1}] &= a \cdot \Delta \mathbb{F}_t[x_{1,t+1}] + b \cdot \Delta \mathbb{F}_t[x_{2,t+1}].\end{aligned}$$

The Internet Appendices A.3 and A.4 contain all proofs.

Next, define the forecast error for a generic variable, η , as realization minus forecast, $\mathbb{F}\mathbb{E}_t[\eta_{t+1}] = \eta_{t+1} - \mathbb{F}_t[\eta_{t+1}]$, and recall that these forecast errors are defined at the CFO (and, hence, firm) level. We obtain the following key result.

Lemma 2 (Coherence, accuracy, and forecast errors under Cobb-Douglas). *Forecast coherence requires that,*

$$\mathbb{F}\mathbb{E}_t[\Delta y_{t+1}] \equiv \mathbb{F}\mathbb{E}_t[\log y_{t+1}] = a \cdot \mathbb{F}\mathbb{E}_t[\log x_{1,t+1}] + b \cdot \mathbb{F}\mathbb{E}_t[\log x_{2,t+1}].$$

Lemma 2 links coherence and accuracy: accuracy requires forecast errors to be zero on average item by item, whereas coherence requires forecast errors of different items to be cross-sectionally linked by the technology parameters. While the exact equality relies on Cobb-Douglas log-linearity, the distinction between coherence and accuracy does not: under CES, coherence yields local monotonicity and sign restrictions rather than exact equalities (see Internet Appendix A.6).

3.1.2 Optimal forecasts and rules of thumb.

We now consider the forecasting problem when a and b are unknown to the forecaster.

Lemma 3 (Unknown parameters). *If parameters a and b are unknown, a forecaster can estimate them using a linear projection operator, with forecasted variables in logs.*

In what follows, we assume past data on output, inputs, and prices are observed by

CFOs and are informative.¹⁴ A CFO can then use such data to estimate the parameters of interest and provide forecasts. We define the following RoTs for $\Delta\mathbb{F}x_{1,t+1}$:

- (R5) Estimate a multivariate linear projection of $\Delta\bar{x}_{1,t}$ on a constant, $\Delta\bar{y}_t$, and $\Delta\bar{x}_{2,t}$, with coefficients c_0 , c_1 , and c_2 , and analogously for input 2. Then define $\Delta\mathbb{F}x_{1,t+1}^{R5} = c_0 + c_1\Delta\mathbb{F}y_{t+1} + c_2\Delta\mathbb{F}x_{2,t+1}^{R5}$, where $\Delta\mathbb{F}x_{2,t+1}^{R5}$ is the forecast of the other input implied by the same multivariate rule.
- (R3) Estimate a univariate linear projection of $\Delta\bar{x}_{1,t}$ on a constant and $\Delta\bar{y}_t$, with coefficients d_0 and d_1 . Then define $\Delta\mathbb{F}x_{1,t+1}^{R3} = d_0 + d_1\Delta\mathbb{F}y_{t+1}$.
- (R2) Define $\Delta\mathbb{F}x_{1,t+1}^{R2} = \Delta\mathbb{F}y_{t+1}$.
- (R1) Define $\Delta\mathbb{F}x_{1,t+1}^{R1} = \frac{1}{k} \sum_{j=1}^k \Delta\bar{x}_{1,t+1-j}$, where $\Delta\bar{x}_{1,t+1-j}$ is allowed to differ from the cost minimizing growth of input 1, $\Delta x_{1,t}^*$, needed to achieve the sales growth target.

Relative to the empirical definitions in Welch (2017), three simplifications apply here: (R5) contains only technologically relevant inputs; (R4) is indistinguishable from (R3) because there is no industry heterogeneity; and (R2) is a special case of (R3) with $d_0 = 0$ and $d_1 = 1$.

From now on, we maintain the following assumption.

Assumption. *In all previous time periods the firm's actual input spending equals the cost-minimizing level, $\bar{x}_{i,s} = x_{i,s}^*$, $\forall s \leq t$, $i = 1, 2$, needed to attain the sales growth target.*

The CFO has access to an observational data set of past input prices, in which the input prices of the CFO's firm represent a single time series. We proceed to obtain a (partial) ranking of these RoTs based on their expected loss.

Proposition 1 (Ranking RoTs under Cobb-Douglas). *In the Cobb-Douglas case, when the comovement of input prices is sufficiently large, a forecasting rule that conditions on all relevant variables (output and other inputs) coincides with the optimal forecast and minimizes expected squared error. In particular,*

$$\Delta\mathbb{F}x_{1,t+1}^{R5} = \mathbb{E}_t [\Delta x_{1,t+1}^*].$$

¹⁴More formally, we assume all firms have the same technology, all CFOs face the same forecasting and cost minimization problems, and past data are generated by optimizing CFOs.

Hence, (R5) is unbiased and yields weakly lower expected loss than (R3), (R2), or (R1). Moreover, (R3) yields weakly lower expected loss than (R1) and (R2). Finally, (R2) yields weakly lower expected loss than (R1).

Proposition 1 yields a benchmark ranking in which the multivariate rule dominates more restrictive rules. We next ask whether using these RoTs predicts systematic FEs.

Implication 1 (Predictable forecast errors). *When forecasts are generated using incoherent rules, forecast errors become systematically predictable from past realizations. In particular, these rules generate overreaction: high past investment growth leads to overly optimistic forecasts and negative forecast errors ex post.*

More formally, consider the local linear projection,

$$\Delta x_{1,t+1}^* - \Delta \mathbb{F}x_{1,t+1} = m_0 + m_1 \Delta \bar{x}_{1,t}. \quad (6)$$

We show in Internet Appendix A.4 that under the coherent rule (R5), $m_0 = m_1 = 0$, so FEs are not predictable from past realizations. Under (R1) and (R2), instead, $m_1 < 0$: after high CapEx growth, CFOs forecast excessively high future CapEx growth, which then mean reverts and generates negative FEs ex post. Thus incoherent heuristics give rise to overreaction, consistent with Table 6.

3.2 Rules of thumb and underinvestment

We have established that incoherent forecasting rules can generate biased input forecasts relative to the cost-minimizing benchmark. We now show how such forecast bias can translate into real underinvestment when adjustment is costly.

Let $\Delta x_{1,t+1}^*$ denote the cost-minimizing growth rate of input 1 required to achieve the sales target ΔFy_{t+1} . Under rule $(Rk) \in \{(R1), (R2), (R3), (R5)\}$, the CFO issues a forecast $\Delta \mathbb{F}x_{1,t+1}^{Rk}$. Suppose that after the forecast is issued, the firm can adjust input growth, but faces convex adjustment costs. Let realized input growth be,

$$\Delta x_{1,t+1} = \Delta \mathbb{F}x_{1,t+1}^{Rk} + \lambda(\Delta x_{1,t+1}^* - \Delta \mathbb{F}x_{1,t+1}^{Rk}), \quad \lambda \in (0, 1], \quad (7)$$

where λ captures the degree of adjustment. When $\lambda = 1$, the firm fully corrects forecast errors and attains the cost-minimizing level. When $\lambda < 1$, adjustment is partial due to

costs, so deviations between forecasted and optimal growth translate into realized under- or overinvestment. Define the investment gap under rule (Rk) as,

$$G_{t+1}^{Rk} = \Delta x_{1,t+1}^* - \Delta x_{1,t+1}. \quad (8)$$

Substituting the adjustment equation yields,

$$G_{t+1}^{Rk} = (1 - \lambda)(\Delta x_{1,t+1}^* - \Delta \mathbb{F} x_{1,t+1}^{Rk}). \quad (9)$$

Thus, realized underinvestment is proportional to the forecast bias, $\Delta x_{1,t+1}^* - \Delta \mathbb{F} x_{1,t+1}^{Rk}$. When forecasts are unbiased (as under coherent (R5)), $\mathbb{E}[G_{t+1}^{R5}] = 0$. When forecasts are downward biased (as under (R1) or (R2)), the expected investment gap is strictly positive:

$$\mathbb{E}[G_{t+1}^{Rk}] > 0 \quad \text{if} \quad \mathbb{E}[\Delta x_{1,t+1}^* - \Delta \mathbb{F} x_{1,t+1}^{Rk}] > 0. \quad (10)$$

Implication 2 (Underinvestment and RoTs). *Under convex adjustment costs, forecasting rules that generate downward-biased input forecasts produce persistent underinvestment relative to the cost-minimizing benchmark. In particular, if $\mathbb{E}[\Delta x_{1,t+1}^* - \Delta \mathbb{F} x_{1,t+1}^{Rk}] > 0$, then $\mathbb{E}[G_{t+1}^{Rk}] > 0$.*

The intuition is immediate. When the CFO underestimates the input growth required to meet the sales target, the firm must revise plans ex post. If adjustment is costly, revision is partial, and realized investment remains below the optimal level. Because incoherent rules such as (R2), and under some conditions (R1), underpredict CapEx growth relative to the cost-minimizing benchmark, they can yield lower investment than (R5). Implication 2 rationalizes our empirical results in Table 9; see derivations in Internet Appendix A.5.

In sum, coherent forecasts require the forecaster to account for the structure of the firm's problem and the joint behavior of output and inputs. Incoherent RoTs can generate lower performance, overreaction, and underinvestment relative to coherent forecasts, consistent with our empirical results.

3.3 A model of bracketing in corporate forecasts

A natural way to interpret incoherent forecasting rules is through the lens of behavioral bracketing. We provide a formal model of bracketing-based forecasting that builds on Lian (2021) and shows how different informational frictions can rationalize the use of simple forecasting rules. Importantly, the empirical patterns documented above do not depend on this particular microfoundation.

Narrowly bracketed forecasts treat each decision in isolation, whereas broadly bracketed forecasts account for interactions across decisions. We capture this idea by introducing noisy signals and recasting the multidimensional forecasting problem as multiple selves playing an incomplete-information, common-interest game. The CFO “capital self” makes forecasts of capital expenditures’ growth by observing noisy signals of output and labor growth. Conversely, the CFO “labor self” makes forecasts of labor expenditures’ growth by observing noisy signals of output and capital growth. In equilibrium, each self does not observe the others’ signals and therefore forms forecasts with imperfect knowledge of the others’ forecasts.

We obtain the following results (full details and proofs in Internet Appendix A.7).

Proposition 2 (Optimal forecasts with noisy signals). *The optimal forecast of x_1 , given signals $\eta_y = \log y + \epsilon_y$ and $\eta_2 = \log x_2 + \epsilon_2$, with $\epsilon_y \sim \mathcal{N}(\mu_y, s_y^2)$, $\epsilon_2 \sim \mathcal{N}(\mu_2, s_2^2)$ is,*

$$\mathbb{E}[\log x_1 | \eta_y, \eta_2] = \mu_1 + \beta_y (\eta_y - \mu_y) + \beta_2 (\eta_2 - \mu_2),$$

$$\text{where } \beta_y = \frac{a\sigma_1^2}{a^2\sigma_1^2 + b^2\sigma_2^2 + s_y^2 - \frac{b^2\sigma_2^4}{\sigma_2^2 + s_2^2}}, \beta_2 = \frac{ab\sigma_1^2\sigma_2^2}{b^2\sigma_2^4 - (\sigma_2^2 + s_2^2)(a^2\sigma_1^2 + b^2\sigma_2^2 + s_y^2)}.$$

Proposition 2 rationalizes a version of (R5) also as second-best optimal and clarifies that the accuracy of the linear projection depends on the signals’ precision.

Corollary 2.1 (Narrow bracketing). *When $s_y^2, s_2^2 \rightarrow +\infty$, the optimal forecast is $\mathbb{E}[\log x_1 | \eta_y, \eta_2] = \mu_1$.*

When all signals are infinitely noisy, the optimal forecast projects the prior mean of the input being forecasted into the future. Thus, Corollary 2.1 provides conditions under which a narrow-bracketing version of (R1) is second-best optimal.

Corollary 2.2 (Univariate projections). *When $s_2^2 \rightarrow +\infty$ and $0 < s_y^2 < +\infty$, the*

optimal forecast is $\mathbb{E}[\log x_1 | \eta_y, \eta_2] = \mu_1 + \beta_y (\eta_y - \mu_y)$, where $\beta_y = \frac{a\sigma_1^2}{a^2\sigma_1^2 + b^2\sigma_2^2 + s_y^2}$.

When the signal of the second input is infinitely noisy and that of the output is noisy but informative, the optimal forecast is a univariate linear projection, where the intercept is still the prior mean of the first input and the slope is a function of the fundamental uncertainty and precision of the output's signal. Thus, Corollary 2.2 provides conditions under which a univariate-projection version of (R3) is second-best optimal.

4. Related Literature

Our evidence on incoherence in firm plans provides a possible mechanism for understanding resource allocation within the firm (e.g., Maksimovic and Phillips 2002, 2013; Hoberg and Phillips 2025). We discuss here our contributions to the study of survey expectations of firms and to behavioral corporate finance and economics.

4.1 *Survey expectations of firms*

Our paper is related to the recent and growing literature studying top executives' beliefs and forecasts about their firm's variables (Ben-David, Graham, and Harvey 2013; Boutros et al. 2025; Campello, Graham, and Harvey 2010; Campello et al. 2011, 2012). This literature typically studies the accuracy of managerial forecasts, one variable at a time. Bloom et al. (2025) show that forecasting firms' own variables is harder than forecasting the aggregate economy. Gennaioli, Ma, and Shleifer (2016) show that both planned and actual investment are explained by CFOs' expectations of earnings growth. Graham (2022) documents that the sales revenue growth forecast is most important for firms' plans and consequences. We confirm that firms make on average accurate forecasts about sales, but we also show that expectations of other variables—most notably, capital expenditures—are much less accurate, and we link this to incoherence. We add to this literature by documenting wide heterogeneity in the ways managerial forecasts link together multiple related balance-sheet variables, some of which lead to incoherence.

4.2 *Coherence and accuracy sides of rationality in behavioral economics and psychology*

The psychology literature has long recognized that rationality of judgements involves both accuracy ('correspondence') and coherence ('consistency'), and has typically maintained that the two are separate properties (e.g., [Hammond 1996, 2007](#); [Osherson, Shafir, and Smith 1994](#); [Gigerenzer and Gaissmaier 2011](#); [Lee and Zhang 2012](#); [Arkes, Gigerenzer, and Hertwig 2016](#)). This literature has struggled to provide a formal framework or direct evidence to define coherence and disentangle it from accuracy.

Behavioral economics has focused on predicting systematic inaccuracy from specific violations of statistics and probability laws, starting from the coherence principle ([de Finetti 1937](#)), which states that subjective probabilities must conform to the logical and probability axioms posited in [Kolmogorov \(1933\)](#). In a series of famous experiments, [Tversky and Kahneman \(1971, 1974, 1983\)](#) document systematic misconceptions of logic, statistics, and probability theory, including the law of large numbers, the conjunction rule, the law of total probability, and Bayes' theorem. Since then, a number of theories and experiments have shown how such misconceptions predict systematic inaccuracy in specific prediction tasks (e.g., [Wright et al. 1994](#); [Rabin 2002](#); [Benjamin, Rabin, and Raymond 2016](#); [Zhu et al. 2022](#)). [Tversky and Kahneman \(1974\)](#) have famously labeled this research program "heuristics and biases," whereby the use of heuristics generates systematic and predictable forecast errors (see also [Thaler 2018](#)). Some scholars have pointed out how some of these results can also be cast in terms of the coherence-accuracy framework. For instance, [Hammond \(1996\)](#), [Tentori, Crupi, and Russo \(2013\)](#), [Arkes, Gigerenzer, and Hertwig \(2016\)](#), [Jönsson and Shogenji \(2019\)](#), and others discuss how the [Tversky and Kahneman \(1983\)](#) conjunction fallacy can be understood as a violation of coherence with respect to probability laws. In this literature, documenting a form of incoherence in one domain typically serves the purpose of predicting future systematic inaccuracy in another domain, rather than jointly measuring coherence and accuracy within the same forecasting setting.

Our main contribution to this literature is to provide a formal framework, in which

coherence and accuracy are defined with respect to the same forecasting task, allowing us to jointly assess forecast accuracy and coherence. To do so one needs both theory and data. In terms of theory, an economic model provides a benchmark against which to judge coherence. This is similar to the use of probability theory for assessing coherence of probabilistic judgments, but crucially economic theory allows us to nest coherence and accuracy in a setting with optimizing agents. In terms of data, observing both forecasts and realizations—and, thus, FEs—for multiple variables simultaneously allows us to jointly assess accuracy and coherence, and to disentangle the two. Accuracy is assessed by testing whether FEs of each variable are ‘sufficiently’ close to zero. Coherence is assessed by testing whether FEs of different variables are ‘sufficiently’ close to one another. In both cases, the extent of ‘sufficiently’ is pinned down by economic theory.

Disentangling coherence and accuracy is important, since under some conditions coherence can be assessed *ex ante*. The intuition is similar to that of the debiasing program in psychology (e.g., [Fischhoff 1982](#)), with the key difference that in our framework one can use economic theory and regression analysis to determine the *ex ante* optimal rule.

4.3 *Behavioral corporate finance*

[Baker, Ruback, and Wurgler \(2007\)](#) organize the behavioral corporate finance literature under the two themes of irrational investors and irrational managers ([Guenzel and Malmendier \[2020\]](#) also discuss other biased parties). The more developed irrational-investors literature studies how rational managers respond to irrational fluctuations in capital market prices. Our study belongs to the less developed and growing irrational-managers literature, in which, while markets are arbitrage free, managerial behavior can be influenced by psychological biases. So far, this stream of research has mostly focused on how psychological traits, such as optimism, overconfidence, or preference for skewness, can distort managerial expectations about the future and, thus, investment and financing decisions (e.g., [Malmendier and Tate 2005](#); [Gervais, Heaton, and Odean 2011](#); [Schneider and Spalt 2014](#)). We study incoherence—a novel managerial trait—and show that it differs from optimism and overconfidence both in theory and in the data. The empirical distinction remains after controlling directly for measures of

optimism and overconfidence (Ben-David, Graham, and Harvey 2013).

4.4 *Behavioral research on bracketing*

Experimental and behavioral research shows decision-makers tend to narrowly bracket, making interrelated decisions in isolation (e.g., Tversky and Kahneman 1981; Read, Loewenstein, and Rabin 1999; Rabin and Weizsäcker 2009; Ellis and Freeman 2024). Thaler (1985, 2018) and Heath and Soll (1996) propose individuals hold a mental account for each decision, thus failing to capture the interdependence among choice variables implied by their budget constraint and utility function.

Economic theories of narrow bracketing and mental accounting include Barberis, Huang, and Thaler (2006) and Rabin and Weizsäcker (2009) on monetary gambles and Hastings and Shapiro (2013) and Lian (2021) on consumption choices. Our theory is closest to Lian (2021), who endogenizes bracketing by modeling a narrow thinker who makes consumption decisions about individual goods with imperfect knowledge of the others and thus faces difficulties at coordinating different decisions.

We add to this literature by developing the first model of bracketing in a firm setting and providing field evidence consistent with narrow bracketing among CFOs making internal plans. As in Lian (2021), there are no explicit mental budgets, and agents make decisions (forecasts) based on different nonnested information in the sense of Blackwell, that is, neither input i 's signal is more informative than $\neg i$'s signal, nor vice versa. These features differentiate Lian (2021) and us from rational inattention models (e.g., Sims 2003; Mackowiak and Wiederholt 2009; Matějka and McKay 2015; Kőszegi and Matějka 2020), and the Gabaix (2014, 2019) sparsity model. Unlike Lian (2021), we study a production setting where decisions are interconnected via the production technology.

5. Conclusion

In this paper we introduce forecast coherence—internal consistency across forecasts of jointly determined variables—as a benchmark for assessing firms' internal plans. We make three contributions. First, we document that over 80% of CFOs make joint

forecasts of CapEx and sales closely aligned with those implied by the heuristics taught in MBA textbooks. Second, we show that forecasts aligned with restricted heuristics, (R1) and (R2), are systematically associated in the cross-section with predictable reversals in forecast errors, and lower ROA and CapEx than forecasts aligned with a less restricted rule, (R5). Third, we build a model of multidimensional forecasting under a coherence constraint, illustrating how reliance on restricted forecasting rules can generate incoherence and resource misallocation within firms, consistent with our evidence.

Taken together, our findings suggest that, whether CFOs implement them, the planning heuristics taught in MBA textbooks provide a useful framework for interpreting firm plans: they are closely aligned with the forecasts CFOs actually make, and they are systematically related to observable firm outcomes.

We view our results as reflecting both a lack of consensus in MBA textbooks and case studies, and a lack of theory and evidence to distinguish among alternative heuristics. Two takeaways follow. First, simple and widely used rules such as (R1) and (R2) are incoherent and should not be used to forecast CapEx. Second, more generally, not all heuristics are the same: some may be useful, others less so, and they should be evaluated with respect to their intended goals. Doing so typically requires both theory and data.

Our results also raise two broader questions. First, if some heuristics perform poorly, why might CFOs still rely on them? We discuss one possible formalization based on behavioral bracketing, which rationalizes incoherent heuristics as a second-best response to the extent of imprecision of signals about future firm variables. At the same time, our evidence does not identify the economic primitives generating incoherence in the data. Distinguishing among behavioral bracketing and other possible microfoundations is an important task for future work. Second, if some heuristics perform poorly, do firms or markets learn over time to make better plans? Answering this question requires going beyond the cross-sectional variation studied here.

We do not address the broader question of what the best possible forecasting method for internal plans should be, beyond the textbook heuristics studied here. Given the nonlinear nature of the firm's problem, richer data and more flexible forecasting methods, when available, may well yield superior plans.

Code Availability. The replication code is available in the Harvard Dataverse at <https://doi.org/10.7910/DVN/KIDK6R>.

References

- Abel, A., and J. C. Eberly. 1994. A unified model of investment under uncertainty. *American Economic Review* 84:1369–84.
- Arkes, H. R., G. Gigerenzer, and R. Hertwig. 2016. How bad is incoherence? *Decision* 3:20–39.
- Baker, M., R. S. Ruback, and J. Wurgler. 2007. Behavioral corporate finance: A survey. In *Handbook of corporate finance: Empirical corporate finance*, ed. B. E. Eckbo, chap. 4, pp. 145–88. Amsterdam, Netherlands: Elsevier.
- Barberis, N., M. Huang, and R. H. Thaler. 2006. Individual preferences, monetary gambles, and stock market participation: A case for narrow framing. *American Economic Review* 96:1069–90.
- Becker, G. S. 1996. *Accounting for tastes*, 1st ed. Cambridge, MA: Harvard University Press.
- Ben-David, I., J. R. Graham, and C. R. Harvey. 2013. Managerial miscalibration. *Quarterly Journal of Economics* 128:1547–84.
- Benjamin, D. J. 2019. Errors in probabilistic reasoning and judgment biases. In *Handbook of behavioral economics: Applications and foundations 1*, vol. 2, eds. B. D. Bernheim, S. DellaVigna, and D. Laibson, chap. 2, pp. 69–186. Amsterdam, Netherlands: Elsevier.
- Benjamin, D. J., M. Rabin, and C. Raymond. 2016. A model of non-belief in the law of large numbers. *Journal of the European Economic Association* 14:515–44.
- Berndt, E. R. 1976. Reconciling alternative estimates of the elasticity of substitution. *Review of Economics and Statistics* 58:59–68.
- Bloom, N., T. Kawakubo, C. Meng, P. Mizen, R. Riley, T. Senga, and J. V. Reenen. 2025. Do well managed firms make better forecasts? *Review of Economics and Statistics*. Advance Access published December 18, 2025, 10.1162/REST.a.1692.
- Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer. 2020. Overreaction in macroeconomic expectations. *American Economic Review* 110:2748–82.
- Boutros, M., I. Ben-David, J. Graham, C. Harvey, and J. Payne. 2025. The persistence of miscalibration. *Review of Financial Studies*. Advance Access published September 10, 2025, 10.1093/rfs/hhaf070.
- Campello, M., E. Giambona, J. R. Graham, and C. R. Harvey. 2011. Liquidity management and corporate investment during a financial crisis. *Review of Financial Studies* 24:1944–79.
- . 2012. Access to liquidity and corporate investment in Europe during the financial crisis. *Review of Finance* 16:323–46.
- Campello, M., J. R. Graham, and C. R. Harvey. 2010. The real effects of financial constraints: Evidence from a financial crisis. *Journal of Financial Economics* 97:470–87.
- Coibion, O., and Y. Gorodnichenko. 2015. Information rigidity and the expectations

- formation process: A simple framework and new facts. *American Economic Review* 105:2644–78.
- Crossley, T., Y. Gong, T. Stinebrickner, and R. Stinebrickner. 2024. Examining income expectations in the college and early post-college periods: New distributional tests of rational expectations. *Journal of the European Economic Association* 22:2700–2747.
- de Finetti, B. 1937. Foresight: Its logical laws, its subjective sources. In *Breakthroughs in statistics*, eds. S. Kotz, N. J. Johnson, and H. E. Kyburg, pp. 134–74. New York: Springer.
- D'Haultfoeulle, X., C. Gaillac, and A. Maurel. 2021. Rationalizing rational expectations: Characterizations and tests. *Quantitative Economics* 12:817–42.
- Ellis, A., and D. J. Freeman. 2024. Revealing choice bracketing. *American Economic Review* 114:2668–700.
- Fischhoff, B. 1982. Debiasing. In *Judgment under uncertainty: Heuristics and biases*, eds. D. Kahneman, P. Slovic, and A. Tversky, chap. 31, pp. 422–44. Cambridge, UK: Cambridge University Press.
- Fogelin, R. 2003. *Walking the tightrope of reason: The precarious life of a rational animal*. Oxford, UK: Oxford University Press.
- Gabaix, X. 2014. A sparsity-based model of bounded rationality. *Quarterly Journal of Economics* 129:1661–710.
- . 2019. Behavioral inattention. In *Handbook of behavioral economics: Applications and foundations 1*, vol. 2, eds. B. D. Bernheim, S. DellaVigna, and D. Laibson, chap. 4, pp. 261–343. Amsterdam, Netherlands: Elsevier.
- Gennaioli, N., Y. Ma, and A. Shleifer. 2016. Expectations and investment. In *NBER macroeconomics annual*, vol. 30, eds. M. Eichenbaum and J. A. Parker, pp. 379–431. Chicago: University of Chicago Press.
- Gervais, S., J. B. Heaton, and T. Odean. 2011. Overconfidence, compensation contracts, and capital budgeting. *Journal of Finance* 66:1735–77.
- Gigerenzer, G., and W. Gaissmaier. 2011. Heuristic decision making. *Annual Review of Psychology* 62:451–82.
- Giustinelli, P., and M. D. Shapiro. 2024. SeaTE: Subjective ex ante treatment effect of health on retirement. *American Economic Journal: Applied Economics* 16:278–317.
- Graham, J. 2022. Presidential address: Corporate finance and reality. *Journal of Finance* 77:1975–2049.
- Graham, J. R., C. R. Harvey, and M. Puri. 2015. Capital allocation and delegation of decision-making authority within firms. *Journal of Financial Economics* 115:449–70.
- Guenzel, M., and U. Malmendier. 2020. Behavioral corporate finance: Life cycle of a CEO career. In *Oxford research encyclopedia of economics and finance*. Oxford, UK: Oxford University Press.
- Hammond, K. R. 1996. *Human judgment and social policy: Irreducible uncertainty, inevitable error, unavoidable injustice*. Oxford, UK: Oxford University Press.
- . 2007. *Beyond rationality: The search for wisdom in a troubled time*. Oxford, UK: Oxford University Press.
- Hastings, J. S., and J. M. Shapiro. 2013. Fungibility and consumer choice: Evidence from


- commodity price shocks. *Quarterly Journal of Economics* 128:1449–98.
- Hayashi, F. 1982. Tobin's marginal q and average q : A neoclassical interpretation. *Econometrica* 50:213–24.
- Heath, C., and J. B. Soll. 1996. Mental budgeting and consumer decisions. *Journal of Consumer Research* 23:40–52.
- Hoberg, G., and G. Phillips. 2025. Scope, scale, and concentration: The 21st century firm. *Journal of Finance* 80:415–66.
- Holthausen, R. W., and M. E. Zmijewski. 2020. *Corporate valuation: Theory, evidence & practice*, 2nd ed. Cambridge, MA: Cambridge Business Publishers.
- Jönsson, M. L., and T. Shogenji. 2019. A unified account of the conjunction fallacy by coherence. *Synthese* 196:221–37.
- Kőszegi, B., and F. Matějka. 2020. Choice simplification: A theory of mental budgeting and naive diversification. *Quarterly Journal of Economics* 135:1153–207.
- Koller, T., M. Goedhart, and D. Wessels. 2020. *Valuation: Measuring and managing the value of companies*, 7th ed. Hoboken, NJ: John Wiley & Sons.
- Kolmogorov, A. N. 1933. *Grundbegriffe der Wahrscheinlichkeitsrechnung*. English translation by N. Morrison, Foundations of the theory of probability, 1956. New York: Springer.
- Lee, M. D., and S. Zhang. 2012. Evaluating the coherence of take-the-best in structured environments. *Judgment and Decision Making* 7:360–72.
- Lian, C. 2021. A theory of narrow thinking. *Review of Economic Studies* 88:2344–74.
- Lovell, M. C. 1986. Tests of the rational expectations hypothesis. *American Economic Review* 76:110–24.
- Luehrman, T. A., and J. L. Heilprin. 2009. Mercury Athletic Footwear, Inc.: Valuing the opportunity. Brief Case 4051, Harvard Business School.
- Mackowiak, B., and M. Wiederholt. 2009. Optimal sticky prices under rational inattention. *American Economic Review* 99:769–803.
- Maksimovic, V., and G. Phillips. 2002. Do conglomerate firms allocate resources inefficiently across industries? Theory and evidence. *Journal of Finance* 57:721–67.
- . 2013. Conglomerate firms, internal capital markets, and the theory of the firm. *Annual Review of Financial Economics* 5:225–44.
- Malmendier, U., and G. Tate. 2005. CEO overconfidence and corporate investment. *Journal of Finance* 60:2661–700.
- Matějka, F., and A. McKay. 2015. Rational inattention to discrete choices: A new foundation for the multinomial logit model. *American Economic Review* 105:272–98.
- Moysan, G., and M. Senouci. 2016. A note on 2-input neoclassical production functions. *Journal of Mathematical Economics* 67:80–86.
- Muth, J. F. 1961. Rational expectations and the theory of price movements. *Econometrica* 29:315–35.
- Oberfield, E., and D. Raval. 2021. Micro data and macro technology. *Econometrica* 89:703–32.
- Osherson, D., E. Shafir, and E. E. Smith. 1994. Extracting the coherent core of human probability judgement: A research program for cognitive psychology. *Cognition* 50:299–

- Pesaran, H., and M. Weale. 2006. Survey expectations. In *Handbook of economic forecasting*, vol. 1, eds. G. Elliott, C. Granger, and A. Timmermann, chap. 14, pp. 715–76. Amsterdam, Netherlands: Elsevier.
- Posner, R. A. 2014. *Economic analysis of law*. 9th ed. Burlington, MA: Aspen Publishing.
- Rabin, M. 2002. Inference by believers in the law of small numbers. *Quarterly Journal of Economics* 117:775–816.
- Rabin, M., and G. Weizsäcker. 2009. Narrow bracketing and dominated choices. *American Economic Review* 99:1508–43.
- Read, D., G. Loewenstein, and M. Rabin. 1999. Choice bracketing. *Journal of Risk and Uncertainty* 19:171–97.
- Ruback, R. S. 2004. Know your worth: Critical valuation errors to avoid. Transcript, Harvard University.
- Schneider, C., and O. Spalt. 2014. Conglomerate investment, skewness, and the CEO long-shot bias. *Journal of Finance* 71:635–71.
- Sen, A. 1993. Internal consistency of choice. *Econometrica* 61:495–521.
- Sims, C. A. 2003. Implications of rational inattention. *Journal of Monetary Economics* 50:665–90.
- Stafford, E., and J. L. Heilprin. 2011. Valuation of AirThread Connections. Brief Case 4263, Harvard Business School.
- Syverson, C. 2011. What determines productivity? *Journal of Economic Literature* 49:326–65.
- Tentori, K., V. Crupi, and S. Russo. 2013. On the determinants of the conjunction fallacy: Probability versus inductive confirmation. *Journal of Experimental Psychology: General* 142:235–55.
- Thaler, R. H. 1985. Mental accounting and consumer choice. *Management Science* 4:199–214.
- . 2018. From cashews to nudges: The evolution of behavioral economics. *American Economic Review* 108:1265–87.
- Titman, S., and J. D. Martin. 2016. *Valuation: The art and science of corporate investment decisions*. 3rd ed. New York: Pearson.
- Tversky, A., and D. Kahneman. 1971. Belief in the law of small numbers. *Psychological Bulletin* 76:105–10.
- . 1974. Judgment under uncertainty: Heuristics and biases. *Science* 185:1124–31.
- . 1981. The framing of decisions and the psychology of choice. *Science* 211:453–58.
- . 1983. Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review* 90:293–315.
- Welch, I. 2017. *Corporate finance*, 4th ed. IAW.
- Wright, G., G. Rowe, F. Bolger, and J. Gammack. 1994. Coherence, calibration and expertise in judgmental probability forecasting. *Organizational Behavior and Human Decision Processes* 57:1–25.
- Zhu, J.-Q., P. W. Newall, J. Sundh, N. Chater, and A. N. Sanborn. 2022. Clarifying the relationship between coherence and accuracy in probability judgments. *Cognition*

223:105022.


ACCEPTED MANUSCRIPT

Figure 1. Coherence restrictions on contemporaneous forecast errors




The figure depicts the joint space of forecast errors (FEs) of sales growth and capital expenditures (CapEx) growth, each computed as realization minus forecast. The shaded areas depict the coherence regions, in which FEs are positively associated with one another and they both lie between the 45-degree line and the horizontal axis. **Alt text:** Two shaded wedges in the first and third quadrants represent coherent combinations of sales and capital-expenditure forecast errors, bounded by the horizontal axis and the 45-degree line.

Figure 2. Contemporaneous forecast errors of output and capital



The figure plots the joint distribution of the forecast errors (FEs) of sales growth and CapEx growth in the data in the same space described in Figure 1. **Alt text:** Scatter plot of sales and capital-expenditure forecast errors. The errors are weakly positively related, with many observations lying outside the coherence regions defined in Figure 1.

Figure 3. (In)coherence and (in)accuracy areas



The figure illustrates the theoretical connection between forecast accuracy and forecast coherence for a hypothetical firm whose capital input has a loading of one-third. There are four conceptual areas, described in the figure. **Alt text:** Conceptual diagram crossing an accuracy region centered on zero forecast errors with a diagonal coherence region implied by a capital loading of one-third, yielding four labeled combinations.

Table 1. Compustat regressions for rules of thumb

The table reports results from ordinary least squares (OLS) estimations. The dependent variable is CapEx Growth, the annual percent growth rate of capital expenditures. The regressors are Past CapEx Growth, the average of the past two lagged annual percent growth rates of capital expenditures; Sales Growth, the annual percent growth rate of total revenues; and Earnings Growth, the annual percent growth rate of earnings. In the first two rows and the last row the sample is the Compustat universe over 1999-2019. In the other rows the sample is the one-digit SIC industry indicated in the first column. CapEx Growth, Sales Growth, and Earnings Growth are winsorized at the 1st and 99th percentiles. Standard errors, two-way clustered by firm and year, are reported in parentheses. * $p < .1$; ** $p < .05$; *** $p < .01$ (two-tailed).

	Intercept	Past CapEx growth	Sales growth	Earnings growth	R^2	No. obs.
Rule 1 ('narrow bracketing')	0.339*** (0.039)	-0.101*** (0.013)			.005	84,119
Rule 3 ('economies of scale')	0.230*** (0.026)		1.080*** (0.034)		.074	110,544
Rule 4 ('industry based')						
SIC 0	0.327** (0.152)		2.037*** (0.381)		.093	370
SIC 1	0.243*** (0.045)		0.972*** (0.074)		.113	8,986
SIC 2	0.245*** (0.026)		0.881*** (0.054)		.050	14,794
SIC 3	0.186*** (0.024)		1.211*** (0.059)		.104	24,880
SIC 4	0.182*** (0.022)		0.961*** (0.095)		.065	14,425
SIC 5	0.163*** (0.027)		1.323*** (0.115)		.082	10,273
SIC 6	0.369*** (0.033)		1.044*** (0.050)		.034	17,358
SIC 7	0.204*** (0.038)		1.186*** (0.091)		.102	14,772
SIC 8	0.197*** (0.023)		1.240*** (0.131)		.087	3,916
SIC 9	0.220*** (0.057)		1.306*** (0.190)		.119	770
Rule 5	0.229*** (0.025)		1.067*** (0.035)	0.017*** (0.004)	.075	109,842

Table 2. Minimum distance of sales forecasts from rules of thumb

The top part of panel A reports selected statistics on the orthogonal distance between the CFO forecast of sales growth and the forecast of sales growth implied by the five rules of thumb. The first column reports distance statistics for the CFOs whose forecast is closest to the forecast implied by R1. The second column reports distance statistics for the CFOs whose forecast is closest to the forecast implied by R2. The third column reports distance statistics for the CFOs whose forecast is closest to the forecast implied by R3. The fourth column reports distance statistics for the CFOs whose forecast is closest to the forecast implied by R4. The fifth column reports statistics for the CFOs whose forecast is closest to the forecast implied by R5. The last column reports statistics for the full sample, whereby each CFO's distance is the minimum distance between that CFO's forecast and the forecasts implied by the five rules of thumb. The bottom part of panel A reports the absolute and relative frequencies of CFOs whose forecast of sales growth is closest to the corresponding rule of thumb by column. Panel B reports the same statistics, assigning to a residual category, R6, the CFOs whose actual forecasts are outside of a ± 0.05 interval around the forecasts implied by R1, R2, R3, R4, and R5.

<i>A. Main categorization R1-R5</i>							
	R1	R2	R3	R4	R5	All	
Mean	0.058	0.030	0.014	0.031	0.044	0.033	
Std. dev.	0.100	0.065	0.017	0.034	0.067	0.059	
Frac. zeros	0.000	0.273	0.000	0.009	0.000	0.109	
P10	0.008	0.000	0.002	0.002	0.004	0.000	
P25	0.015	0.000	0.004	0.007	0.010	0.006	
P50	0.028	0.014	0.007	0.023	0.026	0.018	
P75	0.064	0.035	0.020	0.045	0.044	0.038	
P90	0.114	0.071	0.047	0.068	0.091	0.071	
No. of observations	30	154	39	113	60	396	
Fraction	.076	.389	.098	.285	.152	1.000	
<i>B. Alternative categorization allowing for residual category R6</i>							
	R1	R2	R3	R4	R5	R6	All
Mean	0.020	0.014	0.012	0.018	0.022	0.109	0.033
Std. dev.	0.012	0.012	0.013	0.015	0.015	0.107	0.059
Frac. zeros	0.000	0.326	0.000	0.011	0.000	0.000	0.109
P10	0.006	0.000	0.002	0.001	0.004	0.050	0.000
P25	0.009	0.000	0.004	0.004	0.009	0.064	0.006
P50	0.018	0.014	0.007	0.013	0.020	0.079	0.018
P75	0.028	0.021	0.017	0.028	0.036	0.106	0.038
P90	0.037	0.035	0.034	0.041	0.043	0.143	0.071
No. of observations	20	129	37	89	49	72	396
Fraction (of R1-R6)	.050	.326	.093	.225	.124	.182	1.000
Fraction (of R1-R5)	.062	.398	.114	.275	.151	–	1.000

Table 3. Incoherence and rules of thumb indicators

The top part of the table reports results from OLS estimations. The dependent variable is our measure of incoherence, the minimum orthogonal distance between the CFO forecast of sales growth and the forecast of sales growth implied by Rule 5. The regressors are indicator variables. Rule 1 ('narrow bracketing') equals one if the CFO's sales growth forecast is closest to the forecast implied by R1. Rule 2 ('sales anchoring') equals one if the CFO's sales growth forecast is closest to the forecast implied by R2. Rule 3 ('economies of scale') equals one if the CFO's sales growth forecast is closest to the forecast implied by R3. Rule 4 ('industry based') equals one if the CFO's sales growth forecast is closest to the forecast implied by R4. Rule 5 is the excluded category. OLS standard errors are reported in parentheses. The bottom part of the table reports selected summary statistics of the dependent variable. * $p < .1$; ** $p < .05$; *** $p < .01$ (two-tailed).

	(1)	(2)	(3)	(4)	(5)
Rule 1 ('narrow bracketing')	0.082*** (0.014)				0.103*** (0.016)
Rule 2 ('sales anchoring')		0.039*** (0.008)			0.051*** (0.011)
Rule 3 ('economies of scale')			-0.058*** (0.013)		-0.025* (0.015)
Rule 4 ('industry based')				-0.028*** (0.009)	0.007 (0.011)
Constant	0.065*** (0.004)	0.056*** (0.005)	0.077*** (0.004)	0.080*** (0.005)	0.044*** (0.009)
R^2	.075	.058	.048	.026	.185
No observations	396	396	396	396	396
Summary statistics of the dependent variable					
Mean	0.072				
Std. dev.	0.079				
P10	0.011				
Median	0.057				
P90	0.139				

Table 4. Incoherence and CFO & firm characteristics

Panel A reports descriptive statistics on selected CFO and firm characteristics. Panel B reports results from OLS estimations showing how our measure of incoherence relates to CFO and firm characteristics. The dependent variable is our measure of incoherence, the minimum orthogonal distance between the CFO forecast of sales growth and the forecast of sales growth implied by Rule 5. CFO has MBA is an indicator equal to one if the CFO has an MBA. Age is the age of CFOs expressed in years. Age 40- is an indicator equal to one if age is lower than 40. Age 41-50 is an indicator equal to one if age is between 41 and 50. Age 51-60 is an indicator equal to one if age is between 51 and 60. Tenure is the number of years the CFO has been with their current firm. Tenure > Median is an indicator equal to one if the CFO's tenure is above the sample median of 3. Gender is an indicator equal to one if the CFO is female. Miscalibration is the negative imputed individual volatility, computed as the mean of the imputed individual volatility within each survey date subtracted from the imputed individual volatility; the result is then multiplied by -1 and divided the standard deviation of the imputed individual volatilities within the survey date. Optimism is a standardized expected return measure, computed by subtracting the mean within each survey date from the expected return estimate and dividing by the standard deviation within the survey date. Both miscalibration and optimism apply the same procedure to Short-Term (1-year) and Long-Term (10-year) forecasts. Firm size is the logarithm of total assets. Market-to-book is the book value of assets plus the market value of common equity minus the book value of common equity and deferred taxes divided by total assets. Dividends is an indicator equal to one if common dividend is strictly larger than zero. One-digit SIC industry fixed effects and survey fixed effects are included but their coefficients are not reported. Heteroscedasticity-robust standard errors are reported in parentheses. * $p < .1$; ** $p < .05$; *** $p < .01$ (two-tailed).

A. Summary statistics

	Mean	Std. dev.	Q10	Median	Q90	No. obs.
CFO has MBA	0.452	0.498	0.000	0.000	1.000	396
Age	50.43	6.937	42.00	50.00	60.00	396
Tenure	4.273	4.099	0.000	3.000	9.000	396
Gender	0.088	0.284	0.000	0.000	0.000	396
Miscalibration ST	0.036	0.919	-1.166	0.329	0.985	360
Optimism ST	0.052	0.980	-0.914	-0.077	1.285	373
Miscalibration LT	0.039	0.978	-1.095	0.262	0.917	362
Optimism LT	0.033	1.087	-1.008	-0.078	1.077	374
Firm size	7.829	2.296	4.898	7.805	10.61	396
Market-to-book	1.685	0.900	0.998	1.393	2.731	364
Dividends	0.636	0.482	0.000	1.000	1.000	396

Table 4. Incoherence and CFO & firm characteristics (Continued)

B. Incoherence and CFO characteristics

	(1)	(2)	(3)	(4)	(5)	(6)
CFO has MBA	0.005 (0.009)					0.007 (0.011)
Tenure > Median	0.008 (0.008)					0.006 (0.012)
Age 40-		-0.011 (0.022)				-0.024 (0.029)
Age 41-50		-0.026* (0.015)				-0.038** (0.019)
Age 51-60		-0.023 (0.017)				-0.030* (0.017)
Gender		0.002 (0.010)				0.005 (0.012)
Miscalibration ST			-0.012 (0.008)			-0.014 (0.010)
Optimism ST			-0.012 (0.008)			-0.010 (0.009)
Miscalibration LT				-0.006 (0.004)		-0.002 (0.006)
Optimism LT				0.001 (0.004)		0.005 (0.007)
Firm size					-0.006* (0.003)	-0.005* (0.003)
Market-to-book					0.011 (0.013)	0.015 (0.015)
Dividends					-0.016 (0.012)	-0.020 (0.013)
Constant	0.042* (0.022)	0.077*** (0.025)	0.053* (0.028)	0.045** (0.021)	0.138*** (0.036)	0.159*** (0.049)
Industry FE	Yes	Yes	Yes	Yes	Yes	Yes
Survey FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	.109	.114	.152	.125	.162	.227
No. of observations	396	396	360	362	364	332

Table 5. Regression tests of accuracy and coherence

The table reports results from OLS estimations. The dependent variable in column 1 is the forecast error of log capex growth. The dependent variable in columns 2, 3, 4, and 5 is the forecast error of log sales growth. The forecast error of log capex (re. sales) growth is constructed using the variance-adjusted approximation, $\mathbb{E}_t \log x_{t+1} \approx \log \mathbb{E}_t x_{t+1} - \frac{1}{2} V_t \log x_{t+1}$. In turn, the conditional variance is computed from AR(1) models under stationarity of capital input (re. output) growth using BEA-BLS data and assigned to firms via an SIC-to-BEA industry concordance. Hence, the dependent variable equals the realized log growth minus this variance-adjusted forecast. The forecast error of log wages growth is computed using a similar procedure. Nonparametric bootstrap standard errors clustered at the firm level are reported in parentheses. *p < .1; **p < .05; ***p < .01 (two-tailed).

	FE of log CapEx growth (1)		FE of log Sales growth (2)	(3)	(4)	(5)
FE of log CapEx growth				0.112* (0.063)	0.135*** (0.032)	
FE of log Wages growth				0.023 (0.309)		0.019 (0.321)
Constant	-0.042* (0.025)		-0.009 (0.009)	0.046** (0.023)	-0.004 (0.009)	0.033 (0.022)
R^2	.000		.000	.108	.127	.018
No. of observations	359		360	51	359	52

Table 6. Rules of thumb and predictable forecast errors

The table reports results from OLS estimations. The dependent variable is the forecast error of log capex growth, constructed using the variance-adjusted approximation, $\mathbb{E}_t \log x_{t+1} \approx \log \mathbb{E}_t x_{t+1} - \frac{1}{2} V_t \log x_{t+1}$. In turn, the conditional variance is computed from AR(1) models under stationarity of capital input growth using BEA-BLS data and assigned to firms via an SIC-to-BEA industry concordance. Hence, the dependent variable equals realized log capex growth minus the corresponding variance-adjusted forecast. The main independent variables are Capex Growth, the annual percent growth rate of capital expenditures; Sales Growth, the annual percent growth rate of total revenues; Earnings Growth, the annual percent growth rate of earnings; and four indicator variables. Rule 1 ('narrow bracketing') equals one if the CFO's sales growth forecast is closest to the forecast implied by R1. Rule 2 ('sales anchoring') equals one if the CFO's sales growth forecast is closest to the forecast implied by R2. Rule 3 ('economies of scale') equals one if the CFO's sales growth forecast is closest to the forecast implied by R3. Rule 4 ('industry based') equals one if the CFO's sales growth forecast is closest to the forecast implied by R4. Rule 5 is the excluded category. Nonparametric bootstrap standard errors clustered at the firm level are reported in parentheses. *p < .1; **p < .05; ***p < .01 (two-tailed).

	(1)	(2)	(3)	(4)
CapEx growth	-0.202*** (0.075)	-0.213*** (0.069)	0.023 (0.113)	0.023 (0.110)
Sales growth		0.320 (0.219)		0.298 (0.221)
Earnings growth		-0.056 (0.048)		-0.051 (0.049)
CapEx growth * Rule 1			-0.300 (0.328)	-0.328 (0.315)
CapEx growth * Rule 2			-0.355** (0.147)	-0.338** (0.139)
CapEx growth * Rule 3			-0.075 (0.159)	-0.111 (0.168)
CapEx growth * Rule 4			-0.316 (0.223)	-0.331* (0.195)
Rule 1			0.044 (0.106)	0.063 (0.105)
Rule 2			0.050 (0.068)	0.071 (0.072)
Rule 3			0.123 (0.090)	0.112 (0.089)
Rule 4			0.024 (0.077)	0.032 (0.081)
Constant	-0.003 (0.023)	-0.025 (0.028)	-0.048 (0.062)	-0.079 (0.066)
R^2	.090	.123	.143	.171
No. of observations	359	358	359	358

Table 7. Incoherence and corporate performance (return on assets)

The table reports results from OLS estimations. The dependent variable is return on assets, income before extraordinary items divided by beginning-of-period total assets. The main regressor is our measure of incoherence, the minimum orthogonal distance between the CFO forecast of sales growth and the forecast of sales growth implied by Rule 5. Miscalibration is the negative imputed individual volatility, computed as the mean of the imputed individual volatility within each survey date subtracted from the imputed individual volatility; the result is then multiplied by -1 and divided the standard deviation of the imputed individual volatilities within the survey date. Optimism is a standardized expected return measure, computed by subtracting the mean within each survey date from the expected return estimate and dividing by the standard deviation within the survey date. Both miscalibration and optimism apply the same procedure to Short-Term (1-year) and Long-Term (10-year) forecasts. Firm size is the logarithm of total assets. Market-to-book is the book value of assets plus the market value of common equity minus the book value of common equity and deferred taxes divided by total assets. Dividends is an indicator equal to one if common dividend is strictly larger than zero. One-digit SIC industry fixed effects and survey fixed effects are included where indicated but their coefficients are not reported. Nonparametric bootstrap standard errors clustered at the firm level are reported in parentheses. * $p < .1$; ** $p < .05$; *** $p < .01$ (two-tailed).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Incoherence	-0.381** (0.158)	-0.383** (0.180)	-0.365** (0.162)	-0.398** (0.162)	-0.401** (0.187)	-0.388** (0.170)	-0.318* (0.191)	-0.309* (0.181)
Miscalibration ST		0.003 (0.005)			0.001 (0.005)		-0.001 (0.004)	
Optimism ST		0.000 (0.006)			0.000 (0.006)		0.001 (0.005)	
Miscalibration LT			0.004 (0.005)			0.002 (0.005)		0.001 (0.005)
Optimism LT			0.008 (0.006)			0.007 (0.006)		0.009* (0.006)
Firm size							0.009*** (0.003)	0.009** (0.003)
Market-to-book							0.028** (0.014)	0.028* (0.015)
Dividends							0.022* (0.012)	0.023* (0.013)
Constant	0.069*** (0.010)	0.069*** (0.011)	0.068*** (0.010)	0.053*** (0.013)	0.056*** (0.020)	0.056*** (0.019)	-0.131*** (0.047)	-0.123*** (0.047)
Industry FE	No	No	No	Yes	Yes	Yes	Yes	Yes
Survey FE	No	No	No	Yes	Yes	Yes	Yes	Yes
R^2	.049	.050	.055	.113	.115	.119	.231	.238
No. of CFOs	311	282	284	311	282	284	263	265
No. of firms	277	252	254	277	252	254	235	237
No. of observations	468	423	428	468	423	428	396	401

Table 8. Rules of thumb and corporate performance (Return on Assets)

The table reports results from OLS estimations. The dependent variable is return on assets, income before extraordinary items divided by beginning-of-period total assets. The main regressors are four indicator variables. Rule 1 ('narrow bracketing') equals one if the CFO's sales growth forecast is closest to the forecast implied by R1. Rule 2 ('sales anchoring') equals one if the CFO's sales growth forecast is closest to the forecast implied by R2. Rule 3 ('economies of scale') equals one if the CFO's sales growth forecast is closest to the forecast implied by R3. Rule 4 ('industry based') equals one if the CFO's sales growth forecast is closest to the forecast implied by R4. Rule 5 is the excluded category. Miscalibration is the negative imputed individual volatility, computed as the mean of the imputed individual volatility within each survey date subtracted from the imputed individual volatility; the result is then multiplied by -1 and divided the standard deviation of the imputed individual volatilities within the survey date. Optimism is a standardized expected return measure, computed by subtracting the mean within each survey date from the expected return estimate and dividing by the standard deviation within the survey date. Both miscalibration and optimism apply the same procedure to Short-Term (1-year) and Long-Term (10-year) forecasts. Firm size is the logarithm of total assets. Market-to-book is the book value of assets plus the market value of common equity minus the book value of common equity and deferred taxes divided by total assets. Dividends is an indicator equal to one if common dividend is strictly larger than zero. One-digit SIC industry fixed effects and survey fixed effects are included where indicated but their coefficients are not reported. Nonparametric bootstrap standard errors clustered at the firm level are reported in parentheses. *p < .1; **p < .05; ***p < .01 (two-tailed).

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Rule 1 ('narrow bracketing')	-0.059*** (0.022)	-0.062** (0.025)	-0.061** (0.024)	-0.052** (0.023)	-0.060** (0.025)	-0.056** (0.025)	-0.055** (0.026)	-0.053** (0.026)
Rule 2 ('sales anchoring')	-0.027** (0.014)	-0.029** (0.014)	-0.025* (0.015)	-0.024 (0.015)	-0.029* (0.017)	-0.025 (0.016)	-0.036* (0.022)	-0.033* (0.019)
Rule 3 ('economies of scale')	-0.024 (0.017)	-0.029 (0.019)	-0.027 (0.019)	-0.020 (0.019)	-0.029 (0.021)	-0.026 (0.021)	-0.044* (0.023)	-0.041* (0.022)
Rule 4 ('industry based')	-0.018 (0.012)	-0.017 (0.013)	-0.017 (0.014)	-0.013 (0.013)	-0.015 (0.014)	-0.013 (0.014)	-0.019 (0.015)	-0.018 (0.015)
Miscalibration ST		0.001 (0.005)			-0.001 (0.006)		-0.002 (0.004)	
Optimism ST		0.001 (0.006)			0.001 (0.005)		0.001 (0.005)	
Miscalibration LT			0.004 (0.006)			0.002 (0.005)		0.001 (0.004)
Optimism LT			0.008 (0.006)			0.006 (0.006)		0.008 (0.005)
Firm size							0.010** (0.004)	0.009*** (0.004)
Market-to-book							0.029** (0.013)	0.028* (0.015)
Dividends							0.028** (0.014)	0.029** (0.014)
Constant	0.066*** (0.011)	0.068*** (0.012)	0.066*** (0.013)	0.042*** (0.015)	0.047*** (0.018)	0.048* (0.026)	-0.145*** (0.047)	-0.136*** (0.049)
Industry FE	No	No	No	Yes	Yes	Yes	Yes	Yes
Survey FE	No	No	No	Yes	Yes	Yes	Yes	Yes
R ²	.021	.025	.030	.081	.087	.090	.223	.228
No. of CFOs	311	282	284	311	282	284	263	265
No. of firms	277	252	254	277	252	254	235	237
No. of observations	468	423	428	468	423	428	396	401

Table 9. Rules of thumb and corporate investment

The table reports results from OLS estimations. The dependent variable is investment, capital expenditures divided by beginning-of-period assets. The main regressors are four indicator variables. Rule 1 ('narrow bracketing') equals one if the CFO's sales growth forecast is closest to the forecast implied by R1. Rule 2 ('sales anchoring') equals one if the CFO's sales growth forecast is closest to the forecast implied by R2. Rule 3 ('economies of scale') equals one if the CFO's sales growth forecast is closest to the forecast implied by R3. Rule 4 ('industry based') equals one if the CFO's sales growth forecast is closest to the forecast implied by R4. Rule 5 is the excluded category. Miscalibration is the negative imputed individual volatility, computed as the mean of the imputed individual volatility within each survey date subtracted from the imputed individual volatility; the result is then multiplied by -1 and divided the standard deviation of the imputed individual volatilities within the survey date. Optimism is a standardized expected return measure, computed by subtracting the mean within each survey date from the expected return estimate and dividing by the standard deviation within the survey date. Both miscalibration and optimism apply the same procedure to Short-Term (1-year) and Long-Term (10-year) forecasts. Firm size is the logarithm of total assets. Market-to-book is the book value of assets plus the market value of common equity minus the book value of common equity and deferred taxes divided by total assets. Dividends is an indicator equal to one if common dividend is strictly larger than zero. One-digit SIC industry fixed effects and survey fixed effects are included where indicated but their coefficients are not reported. Nonparametric bootstrap standard errors clustered at the firm level are reported in parentheses. *p < .1; **p < .05; ***p < .01 (two-tailed).

	(1)	(2)	(3)
Rule 1 ('narrow bracketing')	-0.016 (0.011)	-0.014 (0.011)	-0.015 (0.011)
Rule 2 ('sales anchoring')	-0.013** (0.006)	-0.015** (0.007)	-0.012 (0.007)
Rule 3 ('economies of scale')	-0.008 (0.008)	-0.012 (0.009)	-0.011 (0.009)
Rule 4 ('industry based')	-0.002 (0.007)	-0.003 (0.008)	-0.003 (0.008)
Miscalibration ST		0.001 (0.003)	
Optimism ST		0.002 (0.003)	
Miscalibration LT			0.002 (0.002)
Optimism LT			0.004* (0.002)
Firm size	-0.002 (0.002)	-0.002 (0.002)	-0.002 (0.001)
Market-to-book	0.006** (0.003)	0.006** (0.003)	0.006* (0.003)
Dividends	0.000 (0.009)	0.003 (0.008)	0.004 (0.008)
Constant	0.044* (0.024)	0.042* (0.022)	0.049** (0.023)
Industry FE	Yes	Yes	Yes
Survey FE	Yes	Yes	Yes
R^2	.259	.280	.286
No. of observations	437	397	402

Table 10. Change in corporate performance and investment when new CFOs take office

The table reports results from OLS estimations. The dependent variable in columns 1 and 2 is change in ROA, and in columns 3 and 4 is change in investment. ROA is return on assets, income before extraordinary items divided by beginning-of-period total assets. Investment is capital expenditures divided by beginning-of-period assets. The changes in the dependent variables are computed as the value of each variable in the year after the CFO's taking office minus the value in the year before the CFO's taking office. The main regressors are incoherence and four indicator variables. Incoherence is the minimum orthogonal distance between the CFO forecast of sales growth and the forecast of sales growth implied by Rule 5. Rule 1 ('narrow bracketing') equals one if the CFO's sales growth forecast is closest to the forecast implied by R1. Rule 2 ('sales anchoring') equals one if the CFO's sales growth forecast is closest to the forecast implied by R2. Rule 3 ('economies of scale') equals one if the CFO's sales growth forecast is closest to the forecast implied by R3. Rule 4 ('industry based') equals one if the CFO's sales growth forecast is closest to the forecast implied by R4. Rule 5 is the excluded category. Firm characteristics (firm size, market-to-book, and dividends), one-digit SIC industry fixed effects, and CFO-appointment-year fixed effects are included but their coefficients are not reported. Nonparametric bootstrap standard errors clustered at the firm level are reported in parentheses. * $p < .1$; ** $p < .05$; *** $p < .01$ (two-tailed).

	Change in ROA		Change in investment	
	(1)	(2)	(3)	(4)
Incoherence	-1.616*		-0.048	
	(0.982)		(0.045)	
Rule 1 ('narrow bracketing')		-0.268		-0.022*
		(0.213)		(0.012)
Rule 2 ('sales anchoring')		0.000		-0.004
		(0.036)		(0.008)
Rule 3 ('economies of scale')		-0.048		-0.008
		(0.048)		(0.012)
Rule 4 ('industry based')		0.015		-0.000
		(0.046)		(0.009)
Firm characteristics	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes
CFO-appointment-year FE	Yes	Yes	Yes	Yes
R^2	.458	.293	.129	.142
No. of observations	142	142	140	140