The Formation of Very Long-Run Firm Growth Expectations*

Paul H. Décaire† Marius Guenzel‡

August 27, 2026

Abstract

When fundamentals and feedback offer little guidance, do expectations converge on a common anchor or increasingly reflect who forms them? We combine terminal growth rates (TGRs)—expected firm growth in perpetuity—with shorter-horizon forecasts, text, and demographics. TGRs predict long-run growth rather than rationalizing prices. As horizons lengthen, persistent variation shifts from firms to forecasters: firm effects account for 63% of explained variation at one year; forecaster effects account for 61% in terminal beliefs. The forecaster share peaks where common anchors provide least guidance: among young firms and those whose distant growth is least predictable. Forecaster differences reflect experience: analysts extrapolate home-country conditions most strongly to distant but culturally familiar markets. Forecaster-specific variation accounts for 40% of valuation disagreement, 16% of it through terminal growth.

JEL classification: D83, D84, G12, G17, G41.

Keywords: Long-horizon belief formation; terminal growth rates; forecaster heterogeneity; personal experience; firm valuation.

*For their excellent research assistance, we thank Eisa Alloughani, Sebastian Gortarez Ortiz, Tej Kulkarni, Leah Li, Aditeya Pattiwar, Andre Radic, Reed Simons, Cole Streitmatter, Mingyang Zhang, and Shuman Zhang. Décaire gratefully acknowledges financial support from the Center for Responsible Investing at Arizona State University. Guenzel gratefully acknowledges financial support from the Jacobs Levy Equity Management Center for Quantitative Financial Research. Previously circulated under the title: *What Drives Very Long-Run Cash Flow Growth Expectations?* **Researchers interested in accessing the TGR dataset for academic research are encouraged to contact the authors regarding data availability.**

†W. P. Carey School of Business, Arizona State University. Email: paul.decaire@asu.edu.

‡The Wharton School, University of Pennsylvania. Email: mguenzel@wharton.upenn.edu.

1 Introduction

How should expectations about firms behave as the forecast horizon lengthens? One possibility is convergence to a common benchmark. Whatever their near-term views, forecasters may eventually fall back on shared anchors, such as long-run economic growth or mean reversion in profitability, causing individual differences to fade as forecasts approach the terminal horizon. The opposing possibility is that individual differences become more important because the long run is precisely where converging toward a shared benchmark is hardest: both which anchor is relevant and how a given firm approaches it. Current firm dynamics inform near-term forecasts but reveal little about distant growth; the mapping from present signals to long-run fundamentals can drift; and forecasters must wait too long—and observe too few realizations—for feedback to discipline their views. Even when outcomes arrive, forecast errors can confound mistaken beliefs with changes in the underlying process. More generally, when observed outcomes do not reveal how signals map into fundamentals, belief differences grounded in experience and interpretation can persist despite long histories of publicly observable data (Epstein and Schneider, 2007; Acemoglu, Chernozhukov, and Yildiz, 2016). These forces yield opposing predictions: distant expectations may converge toward a common anchor, or they may increasingly reflect the person forming them. The distinction is central to disciplining theories that rely on assumptions about agents' long-run beliefs: some impose convergence to a common benchmark, whereas others take persistent person-specific long-run beliefs as a primitive.[1]

Empirically distinguishing between these views requires repeatedly observing distant beliefs in a setting where *many* forecasters evaluate *many* of the same targets, allowing us to separate persistent differences across the objects being forecast from persistent differences across the forecasters themselves. Firm valuation provides such a setting and makes the distinction economically consequential. To value a firm, forecasters must form expectations about cash flows at both short and very long horizons, and the same firm is typically valued by several forecasters, each covering many firms. Because more than 70 percent of a firm's value typically derives from cash flows arriving beyond the five- to ten-year explicit forecast window (Décaire and Graham, 2026), even small differences in long-run cash-flow beliefs can generate material valuation disagreement. Yet systematic evidence on how market participants form these beliefs remains scarce, and most available data capture only shorter-horizon expectations (Brunnermeier et al., 2021). The expectations carrying the greatest weight in firm valuation are thus the ones we observe least.

[1] Examples of theories assuming convergence toward common long-run anchors across agents include Brock and Hommes (1997); Andrade et al. (2016); Pástor and Veronesi (2003); Bansal and Yaron (2004); Collin-Dufresne, Johannes, and Lochstoer (2016); Ehling, Graniero, and Heyerdahl-Larsen (2018); Gabaix (2020); Nagel and Xu (2022); de Silva and Thesmar (2024); Farmer, Nakamura, and Steinsson (2024); examples of theories in which long-run expectations diverge persistently across agents include Harrison and Kreps (1978); Kogan et al. (2006); Bhamra and Uppal (2014); Borovička (2020).

Our main result is that the source of persistent variation in professional growth forecasts shifts with the horizon: firm differences dominate near-term expectations, whereas forecaster differences become monotonically more important and dominate terminal beliefs—whether measured relative to explained or total variation. This shift does not simply reflect greater disagreement at longer horizons, but a change in what persistent forecast variation represents: the balance shifts progressively from firm differences toward forecaster differences as the horizon lengthens. The expectations that carry the greatest valuation weight are thus not only the ones we observe least; they are also the ones most shaped by who forms them. Across forecast horizons and the firm life cycle, we show that the forecaster component grows as direct information becomes less informative and common anchors weaken. Using analyst–firm distance, we then examine what enters individual judgment under these conditions, connecting the forecaster component to analysts' local economic experiences and to the different weights they place on common macroeconomic, industry, and geographic information. Finally, we show why this shift matters for valuation: because terminal growth assumptions compound in perpetuity, even modest forecaster-specific differences translate into material disagreement in professional valuations.

Together, these exercises point to a broader feature of belief formation: person-specific experience and interpretation matter more when common anchors are weak and direct information is limited. The same forces can shape beliefs wherever fundamentals provide little guidance and feedback offers little discipline, even if their quantitative importance varies across settings.

We study these beliefs using a new dataset of professional forecasters' *terminal growth rate (TGR)* expectations, the growth rates they expect firms to sustain *in perpetuity*. Unlike the widely used IBES long-term growth forecast (LTG), which covers expected growth over the next three to five years, TGRs measure beliefs about growth beyond the explicit forecast period—to our knowledge, the first such measure at the firm level. We obtain 41,000 manually verified TGR forecasts from sell-side equity reports and match each to the same analyst–firm's shorter-horizon IBES forecasts. The reports jointly disclose each analyst's subjective discount rate, terminal horizon, TGR, and price target. The sample covers 7,975 firms in 48 countries between 2000 and 2023 and 5,229 analysts; about one-third of forecasts are issued by analysts outside the firm's country. We also link observations to report text and, for 2,555 analysts in 45 countries, analyst characteristics from Revelio Labs. The rich cross-sectional structure of our panel allows us to distinguish persistent differences across firms from persistent differences across forecasters and to trace their relative importance across horizons.

We begin by establishing that TGRs capture meaningful beliefs about long-run firm growth and measure a different segment of the expected growth term structure than IBES LTG. TGR and IBES LTG forecasts have a firm-year correlation of only 0.14, and TGRs are roughly ten percentage points lower on average, placing them near economy-wide long-run growth rates. More importantly,

TGRs are better calibrated to cross-sectional differences in realized growth at distant horizons than LTG. Higher TGRs predict higher realized growth over horizons extending from three to thirteen years, both unconditionally and after extensive firm and macroeconomic controls. The contrast becomes increasingly pronounced as the realization horizon expands: the association between LTG and realized growth declines sharply, whereas the association for TGR remains economically meaningful. At the longest horizons, the TGR coefficient is roughly fifteen times the LTG coefficient, and once TGR is included, LTG retains virtually no independent association with long-run realized growth. TGRs also vary systematically with firm fundamentals, declining over the firm life cycle and with long-horizon default risk. Together, these patterns indicate that TGRs are not merely generic macroeconomic, risk, or life-cycle assumptions appended to valuation models. They contain firm-specific information about the cross-sectional ranking of long-run growth prospects.

We next address whether TGRs simply rationalize observed prices. If analysts infer growth from prices using discount rates that differ from those of market participants, their reported growth expectations may embed discount-rate wedges rather than independent cash-flow beliefs (Nagel, 2024). Because we jointly observe TGRs and subjective discount rates in analysts' own valuation models, we can test this interpretation directly. Reverse-engineering through TGR requires subjective discount-rate wedges to be offset by corresponding adjustments to terminal growth. Within analyst–firm relationships, after absorbing common macroeconomic variation, we find no such offsetting comovement. Analysts may not reverse-engineer their models, or may instead reconcile them with prices through near-term forecasts; either way, TGR is not simply the residual needed to make the valuation model *fit* observed prices.

We then turn to the paper's central question: how does the source of growth expectations differ across forecast horizons? In the full aggregate sample, persistent firm differences account for 63 percent of the explained variation in one-year-ahead forecasts, compared with 32 percent for forecaster differences.[2] At the terminal horizon, this ranking reverses: firm differences account for 31 percent of explained variation, while forecaster differences account for 61 percent. Applying the same decomposition to LTG forecasts, whose three-to-five-year horizon lies between these endpoints, yields firm and forecaster shares that also fall between them. As the forecast horizon lengthens, the persistent drivers of variation shift gradually from firm characteristics to forecaster characteristics. Measured relative to total variation, the firm share changes comparatively little, from 14 to 21 percent, whereas the forecaster share rises from 7 to 40 percent. Brokerage-house fixed effects explain a negligible share of variation at all three horizons—one year, LTG, and TGR—indicating that the forecaster component arises primarily within, rather than across, brokerage houses.

To establish that this reversal reflects the forecast horizon rather than differences in sample

[2] Shares are averages across the eight specifications of Table 4; the ranking at each horizon holds in every specification.

composition or forecast timing, we repeat the decomposition in a matched sample. We link each TGR to the same analyst's one-year-ahead growth forecast from IBES, directly comparing near-term and terminal expectations within analyst–firm–forecast-date pairs. Although this matching substantially reduces the sample size, it holds fixed the forecasters, firms, and forecast dates. The shift remains large: the forecaster share rises from 38 percent for one-year expectations to 71 percent for terminal expectations, while the firm share falls from 61 to 25 percent. Under a common terminal anchor, differences in learning speeds may generate near-term disagreement, but persistent forecaster differences should be weaker, not stronger, at the terminal horizon. Instead, they become more important, indicating that forecasters differ in their choice of long-run anchor or in how they map individual firms onto it.

Long-horizon disagreement alone cannot establish persistent person-specific anchors: sampling error in estimated fixed effects can inflate conventional variance-component estimates (Kline, Saggio, and Sølvsten, 2020), the more so if expectation noise becomes more important with the forecast horizon (de Silva and Thesmar, 2024). We address this concern by asking whether the same forecaster-specific component carries across independently evaluated firms. Following Chan and Chen (2026), we divide each analyst's firm relationships into two nonoverlapping samples and estimate the effects separately. Because the samples share no firms and the effects are estimated independently, their covariance isolates the part of the analyst effect that carries across firms rather than sampling error or views specific to a particular firm. The effects covary strongly, implying a persistent analyst-effect standard deviation of 0.63 percentage points—about half the cross-sectional standard deviation of TGRs. Applying the same split-sample correction to the one-year and LTG decompositions leaves our main conclusion unchanged: the forecaster component becomes increasingly dominant as the forecast horizon lengthens.

Likewise, a persistent forecaster component could in principle reflect reporting conventions rather than beliefs—round-number defaults, mechanical repeated use of the same terminal assumption, or employer policies. We perform several refinements to rule out these concerns. We first replace brokerage-house fixed effects with brokerage house-by-industry-by-year fixed effects, allowing employer practices to vary flexibly across industries and over time and, in the process, conservatively attributing to employers a substantial share of variation that would otherwise be assigned to firms. The analyst share remains virtually unchanged and significantly exceeds the share attributed to employers. We exclude from the decomposition TGR values associated with round numbers—whole percentage points and half-steps—and find similar magnitudes. Similarly, focusing the analysis on analysts who switch employers only does not affect our conclusions. Overall, our interpretation does not require frequent belief revisions; it requires only that stable forecaster-specific differences carry across firms.

We next examine what drives the strength of the forecaster component. If individual judgment

fills the gap left by weak common anchors, forecaster-specific differences should be largest where a common benchmark has the least to say about distant growth and should decline as informative benchmarks emerge. The firm life cycle provides an intuitive first test. For firms younger than 15 years, country and industry factors account for only 17–18 percent of explained TGR variation, leaving the remainder largely to analyst effects. As firms mature, the share attributable to these factors rises substantially—exceeding 50% for firms 45 years and older—overall TGR dispersion falls, and consequently analyst effects gradually explain less of the variation.

In the life cycle test, firm age proxies for how informative a common benchmark can be about distant growth. We next measure this directly. For each firm-year, we estimate out of sample how well public information—consensus forecasts, fundamentals, and industry and country conditions—predicts realized growth five to ten years ahead and sort firm-years into terciles of this predictability. The same pattern emerges. In the least-predictable tercile, country and industry factors explain less TGR variation, while the analyst share rises to 81%, from 69% in the most-predictable tercile, and the analyst component is about 40% larger in absolute terms.[3] A fixed analyst reporting convention should neither fade as covered firms mature nor receive greater weight where public information predicts poorly, making the joint pattern difficult to attribute to mechanical persistence. The two tests are complementary: firm maturity provides an intuitive cross-sectional counterpart to the horizon result, while the direct predictability test zeroes in on the mechanism underlying the life-cycle pattern—forecaster-specific anchors carry the most weight precisely where common benchmarks have the least to say.

Our second exercise asks what enters individual judgment when common anchors are weak. We exploit variation in the geographic, social, and cultural distance between analysts' home countries and the foreign firms they value. Analysts whose home economies grew faster in recent years assign higher TGRs to foreign firms. Importantly, this relationship reflects two opposing shifts in the weights analysts place on regional economic conditions: as geographic distance increases, the weight placed on growth in the firm's country approaches zero, while the weight placed on growth in the analyst's home country rises significantly. Shared shocks across the two countries cannot by themselves account for this reallocation of weights, and comparing close and distant pairs mitigates the concern that analysts misattribute shocks in neighboring countries to their own. Among geographically distant firms, extrapolation from local experience is strongest when the firm's country is connected to the analyst's home country by strong social ties, a common language, or similar cultural influences. The pattern appears separately across all three measures and persists after

[3] Two unreported checks yield similar results. Measuring predictability using realized growth from $t + 3$ to $t + 13$, rather than $t + 3$ to $t + 8$, increases the analyst share of total TGR variation from 40% to 48% across terciles and the analyst component from 0.60 to 0.79 squared percentage points. Forming terciles within age bracket, country, and industry produces numbers nearly identical to Panel C—an analyst share of 41% to 49% and an analyst component of 0.65 to 0.83—confirming that the predictability sort does not simply relabel firm age.

controlling for bilateral economic ties. This suggests that analysts project their own experience onto distant settings, and most strongly when social or cultural familiarity makes that experience most transferable to the target. Because analysts rarely relocate—fewer than 6% do so in our sample—this tendency to project home-country experience can generate persistent differences across forecasters. Consistent with this interpretation, variation in analyst effects is organized primarily along two dimensions that index experienced macroeconomic history—where an analyst lives and when she was born.[4] Together, these directional patterns give part of the forecaster component economic content: it reflects macroeconomic conditions analysts have experienced rather than residual noise.

Finally, the text of the reports provides corroborating evidence. Analysts assigning different TGRs to the same firm place different weights on the same macroeconomic, industry, and geographic information—differential interpretation in the sense of Kandel and Pearson (1995)—and those weights diverge more between analysts located farther from one another. The textual evidence reinforces the interpretation from the horizon, life-cycle, and geographic results: as common anchors weaken, analysts place greater weight on their own experience and interpretation.

To assess the valuation consequences of this shift in beliefs, we ask how much TGRs contribute to valuation disagreement, and specifically quantify the role of persistent forecaster-specific differences. For each report, we reconstruct a synthetic price-to-earnings ratio from the analyst's own inputs and decompose the disagreement between analysts covering the same firm in the same year. Persistent analyst effects account for about 40 percent of this disagreement. Of that forecaster component, 16 percent flows through terminal growth, about a third through one-year growth, and roughly half through the discount rate. We interpret these exercises as quantifying disagreement in professional valuations, not equilibrium price effects. These magnitudes are large given the underlying differences: analysts covering the same firm differ by only 1 percentage point in terminal growth on average, compared with 11 percentage points in their next-year growth expectations. This points to a material role for forecaster-specific beliefs in shaping valuations: the input most shaped by the forecaster is also the one whose effect compounds in perpetuity. With three-quarters of reconstructed value lying beyond the explicit forecast horizon, a 1-percentage-point difference in terminal growth moves value by roughly 12 percent—an order of magnitude more than the same difference in next-year growth.

To interpret these magnitudes, we adapt the constant-gain learning framework of Delao, Han, and Myers (2026) to a cross-section of forecasters who may differ both in how quickly their expectations converge and in the terminal anchor toward which they converge. Under a common anchor, heterogeneous convergence speeds generate substantial disagreement at intermediate horizons, but

[4] Observable characteristics account for 17–31 percent of the variation in estimated analyst effects, almost entirely through location and birth cohort. Giglio et al. (2021) report R^2s of 2–7 percent from an analogous decomposition for wealthy U.S. retail investors. Prior work identifies characteristics and experiences that predict expectations (e.g., Malmendier and Nagel, 2011; Kuchler and Zafar, 2019; Das, Kuhnen, and Nagel, 2020; D'Acunto et al., 2023), but these observables typically explain only a small share of overall belief heterogeneity.

disagreement vanishes as the horizon lengthens; heterogeneous terminal anchors instead sustain it in the limit. Disciplining terminal-anchor dispersion at 0.355 percentage points, half the median within-firm-year standard deviation of TGRs, heterogeneous anchors alone generate a coefficient of variation in valuations of 4.0 percent, compared with 3.4 percent from considerable heterogeneity in learning speeds around a common anchor. Modest differences in where expectations converge thus generate as much valuation disagreement as large differences in how quickly they converge.

2 Related Literature

Our results contribute to four literatures.

First, we provide evidence on a primitive component that separates theories of expectation formation in economics and finance: whether long-run beliefs converge to a common anchor or contain persistent person-specific determinants (Section 3). Models with common anchors attach the persistent long-run object either to the firm or to a single representative belief: investors may learn about a firm-specific mean, diagnostic distortions may decay toward an unconditional benchmark, and representative agents may learn slowly about low-frequency parameters (Pástor and Veronesi, 2003; Bordalo, Gennaioli, and Shleifer, 2018; Farmer, Nakamura, and Steinsson, 2024). Multi-agent models can nonetheless generate horizon-dependent disagreement through heterogeneous forecasting rules or transient noise around a common, slowly drifting component (Brock and Hommes, 1997; Andrade et al., 2016; de Silva and Thesmar, 2024). None of these mechanisms contains a persistent person-specific terminal component, so long-horizon disagreement alone does not identify one. Other models instead make the subjective long-run law of motion person-specific. Bhamra and Uppal (2014) provide the closest analogue to our object: each agent holds a constant belief about expected growth that enters her valuation of cash-flow claims; related models take persistent differences in subjective probabilities as primitives (Harrison and Kreps, 1978; Kogan et al., 2006; Borovička, 2020). A persistent terminal component and a rising forecaster share are distinct facts: individual anchors permit the former without implying the latter (Martin and Papadimitriou, 2022). Existing evidence documents horizon-dependent disagreement, persistent relative optimism, and common person components across expectations of aggregate outcomes (e.g., Patton and Timmermann, 2010; Crump et al., 2025; Li, Van Nieuwerburgh, and Wang, 2026; Giglio et al., 2021), and, among retail investors, disagreement between short- and long-horizon forecasters about the same stock that predicts trading (Cookson, Dim, and Niessner, 2026), but lacks a crossed forecaster–target panel in which forecaster and target components can be compared and a permanent forecaster component identified across independently evaluated targets. Our many-firms-by-many-forecasters structure provides both comparisons: the forecaster share rises with the horizon, and the same permanent component appears across nonoverlapping sets of firms.

Second, we contribute to research using subjective cash-flow expectations to understand firm valuation by introducing a new direct measure of beliefs about growth beyond the explicit forecast period. Since Porta (1996), this literature has largely relied on IBES LTG—expected growth over the next three to five years—to study valuation ratios, valuation cycles, subsequent returns, and learning from experienced growth (Bordalo et al., 2024; Delao and Myers, 2021, 2024; Nagel and Xu, 2022). Delao, Han, and Myers (2026) combine consensus forecasts through four years with a model in which a representative agent learns about each firm's mean earnings growth, so that long-run beliefs are firm-specific but common across forecasters. We instead observe the terminal growth rate in each analyst–firm valuation and thus measure, rather than infer, expected growth in perpetuity at the forecaster level. Relative to research on valuation practice, discount rates, and analysts' mental models (Décaire and Graham, 2026; Decaire, Sosyura, and Wittry, 2024; Bastianello, Decaire, and Guenzel, 2025a,b), we exploit the analyst-level structure of equity reports to study how terminal beliefs form and enter valuations.

Third, we contribute to research on the origins of persistent belief heterogeneity. Long-run parameters are difficult to learn from finite histories (Farmer, Nakamura, and Steinsson, 2024); personal and local experiences shape beliefs (Malmendier and Nagel, 2011, 2016; Kuchler and Zafar, 2019); and individuals interpret the same information through different mechanisms or representations (Kandel and Pearson, 1995; Kohlhas and Walther, 2021; Andre et al., 2022; Andre, Schirmer, and Wohlfart, 2026). Existing settings rarely combine rich forecaster backgrounds with repeated forecasts of comparable targets. We show that analysts extrapolate home-country conditions most strongly to firms that are geographically distant but culturally familiar, while forecaster dispersion recedes as firms mature and expectations align more closely with country and industry benchmarks. Report text provides complementary evidence that same-firm TGR differences covary with the information analysts emphasize. Our design parallels research using repeated decisions over comparable cases to isolate stable person components (Bertrand and Schoar, 2003; Chan, Gentzkow, and Yu, 2022; Chan and Chen, 2026), but adds a horizon dimension: holding forecasters and firms fixed, the forecaster share grows as fundamentals and feedback provide less guidance.

Finally, we connect persistent heterogeneity in professional expectations to valuations. Prior work documents persistent individual components across investor expectations and links them to portfolio holdings (Giglio et al., 2021), while related work finds forecaster-specific differences in analyst behavior (Antoniou, Frydman, and Kilic, 2025). Neither observes the valuation inputs needed to isolate how a persistent terminal-growth component affects the value assigned to a particular firm. Analysts' disclosed models allow us to test directly whether terminal-growth assumptions absorb subjective discount-rate wedges, the concern that growth expectations are backed out from prices (Nagel, 2024; Chaudhry, 2025), and we find no such relation; this complements prior tests based on realized growth, returns, and expected-return proxies (Bordalo et al., 2024, 2025). Terminal growth

accounts for about 40 percent of within-firm-year valuation disagreement is persistent and forecaster-specific; a sixth of that flows through terminal growth. Extending the representative-forecaster calibration of Delao, Han, and Myers (2026) to multiple forecasters shows that modest differences in terminal anchors generate at least as much valuation disagreement as substantial differences in convergence speeds around a common anchor. Persistent person-specific beliefs therefore materially affect valuations through the cash flows carrying the greatest weight in firm value.

3 A Conceptual Framework

We use a simple conceptual framework to illustrate the paper's central tension: whether heterogeneity at long horizons reflects different paths toward a shared anchor or differences in the anchors themselves. The section's goals is to isolates the primitives on which the two views differ to then organize and motivate our empirical analysis.

Let forecaster j's expectation of firm i's growth approach a long-run anchor geometrically:

$$E_t^j[g_{i,t+s}] = g_{\infty,i,j} + \rho_j^{s-1}\left(E_t^j[g_{i,t+1}] - g_{\infty,i,j}\right), \qquad 0 < \rho_j < 1, \tag{1}$$

where ρ_j governs the speed of convergence.[5]

Equation (1) nests two views that differ in a single primitive: whether the anchor depends on the forecaster,

$$\underbrace{g_{\infty,i,j} = \bar{g}_i}_{\text{common anchor}} \qquad \text{versus} \qquad \underbrace{g_{\infty,i,j} = \bar{g}_i + a_j}_{\text{individual anchors}}. \tag{2}$$

Under a common anchor, forecasters may differ in their information, near-term views, convergence speeds, and the distortions their forecasts carry along the way, but they converge to the same place; the anchor $\bar{g}_i$ may be firm-specific or shared across firms, as with a common long-run growth rate. Under individual anchors, the long-run anchor itself contains a persistent forecaster-specific component, a_j, that the same forecaster carries across the firms she evaluates.[6]

Because ρ_j^{s-1} vanishes as the horizon lengthens, holding the firm and forecast date fixed,

$$\lim_{s\to\infty} \mathrm{Var}_j\left(E_t^j[g_{i,t+s}] \,\middle|\, i,t\right) = \mathrm{Var}_j\left(g_{\infty,i,j} \,\middle|\, i,t\right) = \mathrm{Var}_j\left(a_j \,\middle|\, i,t\right), \tag{3}$$

[5] The restriction $0 < \rho_j < 1$ defines an anchored process. The anchor may vary over calendar time without affecting the distinction below; we suppress this time dependence and focus on its variation across forecasters at a given date.

[6] More generally, the loading on the forecaster-specific component may vary across targets and the mapping may contain a forecaster–firm component, $g_{\infty,i,j} = \bar{g}_i + \omega_{i,t}\, a_j + m_{i,j}$, where $\omega_{i,t}$ captures how much weight the forecaster's own component receives for a given target and $m_{i,j}$ a component specific to the forecaster–firm pair. The limiting variance in equation (3) then becomes $\omega_{i,t}^2\,\mathrm{Var}_j(a_j \mid i,t)$ plus the variance of $m_{i,j}$. We use the additive form only for exposition.

where the final equality follows under the additive representation in equation (2). The limiting variance is zero under a common anchor but can remain positive under individual anchors. Heterogeneous information, near-term beliefs, and convergence speeds can therefore generate disagreement at finite horizons without producing a persistent forecaster-specific component in the terminal belief. Such a component requires the anchor itself, or the forecaster's mapping of targets onto it, to differ persistently across forecasters. The framework identifies heterogeneity in subjective endpoints but is agnostic about its source: persistent priors, experience, reliance on common defaults, differences in the timing of forecasters' information, and reporting practices can all generate it.

The contrast between the two views yields three empirical implications: across horizons, across targets, and across environments. First, at the terminal horizon: under a common anchor, forecaster-specific variation vanishes as the horizon lengthens; under individual anchors, a persistent component remains. Second, across targets: the same forecaster-specific component appears in the terminal beliefs a forecaster holds about different firms, whereas under a common anchor no such component exists to carry across.[7] Third, across environments: if the weight on the forecaster's own component, $\omega_{i,t}$, rises where common benchmarks provide less guidance, and if that component is grounded in the forecaster's own experience, forecaster-specific variation should be largest for targets about which common information is weakest and should bear the imprint of the forecaster's environment. The first implication distinguishes heterogeneity in paths toward a common anchor from heterogeneity in the anchor itself; the second and third characterize when the forecaster-specific component matters and what gives it content.

4 Institutional Details and Data

Since we assemble our dataset of long-run cash-flow expectations directly from the original reports written by professional equity analysts, some institutional background is useful. Analysts rely primarily on two methods to value firms and form 12-month price targets, the main quantitative output of equity reports: discounted cash flow models (DCF) and valuation multiples. Financial textbooks and business school curricula emphasize both approaches (Berk and DeMarzo, 2019; Brealey, Myers, and Allen, 2022), and practitioner surveys confirm their widespread use (Graham and Harvey, 2001; Mukhlynina and Nyborg, 2020; Bastianello, Decaire, and Guenzel, 2025a,b). In the universe of reports housed by our data provider, Refinitiv, one or both appear in nearly every report: 62% use multiples to derive price targets and 43% use DCF (see Section 4.1.1).

The terminal growth rate plays a central role in DCF models.[8] For the modal DCF in our sample,

[7] The common anchor may vary across firms or be shared across them; what matters is that it does not vary across forecasters for a given firm.

[8] Valuation multiples also embed assumptions about terminal growth but do not disclose those assumptions

explicit cash-flow forecasts extend three years, with all subsequent cash flows valued as a growing perpetuity at the TGR. The terminal value consequently accounts for roughly three-quarters of the value underlying the analyst's price target, and a one-percentage-point change in the TGR moves the implied value by approximately 12%—an order of magnitude more than the same change in next-year growth (Section 8). Although other aspects of analysts' broad role may receive greater attention, our analysis requires only that they treat the TGR as an economically meaningful input rather than a mechanical default. The TGR is disclosed in the DCF table used to justify the price target, a salient output on which analysts are evaluated by clients and employers. Consistent with the TGR reflecting deliberate economic judgment, we show below that TGRs track firm fundamentals, predict realized growth at horizons of up to thirteen years, and are not residuals chosen to reconcile the valuation model with observed prices.

Our data collection focuses on reports that use DCF models because they disclose the underlying cash-flow forecasts and valuation assumptions, whereas multiples typically provide a single valuation ratio.[9] Appendix Figure B.1 presents an example of a complete DCF model from an equity report, with its numerical inputs redacted for copyright reasons. This granularity allows us to extract each TGR and distinguish it from the model's other inputs; every terminal growth expectation in our analysis is observed directly in an analyst report rather than estimated or approximated by us. This feature distinguishes our setting from prior studies using commercial databases such as IBES or Value Line, which report only shorter-horizon expectations and therefore do not permit researchers to observe truly long-run beliefs directly.[10]

4.1 Data

4.1.1 TGR Beliefs Data from Analyst Equity Reports

To construct the long-run cash flow expectations dataset, we initially downloaded 157,549 equity reports mentioning "DCF" from 42 major equity research firms via Refinitiv Eikon. We restricted the time window to reports published in the first quarter of each calendar year (January 1 to April 1) from 2000 to 2023, to ensure that belief data and other variables were systematically measured at a

separately.

[9] Analysts do not randomly choose between multiples- and DCF-based methods. Prior work shows that DCF models and cash-flow–based valuation inputs are used more frequently for harder-to-value firms (Asquith, Mikhail, and Au, 2005; Hashim and Strong, 2018). Recent work using the near-universe of analyst reports shows that valuation methods and the reasoning underlying them are systematic and that, conditional on detailed controls, DCF-based valuations are less accurate on average than simpler approaches, although they outperform when used by sufficiently skilled analysts (Bastianello, Decaire, and Guenzel, 2025a,b). Our sample is thus plausibly selected along some dimensions, a tradeoff for directly observing long-run beliefs that multiples-based valuations leave unobserved.

[10] As we discuss in Section 5.1, the "long-term" growth estimates reported in these databases, such as IBES LTG, do not capture analysts' truly long-run expectations: they are approximately an order of magnitude larger than observed TGRs.

consistent time each year. When analysts published multiple reports on the same firm within the first quarter of a given year, we retained only the earliest publication to avoid duplicates within an analyst-firm-year pair. This procedure resulted in 40,819 reports that (1) contained the forecaster's TGR expectation and, crucially for our analysis of heterogeneity in long-run expectations across forecasters, (2) allowed us to unambiguously identify both the analyst's name and location at the time of report publication, using landline country indicators for the latter.

We then extract forecasters' TGR expectations (as well as additional relevant variables such as subjective discount rates) from each report in five steps. First, we pre-process each document using a series of Python programs, including PDF extraction functions and OCR packages, to identify all sections of text, tables, and figures that contain relevant quantitative information for our purposes. Second, we convert these different elements into short text snippets. Third, for each collected variable, we use the Open AI API to extract the numerical value from these snippets. Fourth, to ensure the highest level of accuracy, we export the text snippets and AI-extracted numerical values to Excel and manually verify each extracted value with the help of our research team. Finally, we trim the TGR expectations at the 0.5th and 99.5th percentiles to mitigate the influence of outliers and any residual transcription errors.

4.1.2 Analyst Backgrounds from Universe of LinkedIn Data

Next, we merge the TGR dataset from Section 4.1.1 with the universe of LinkedIn users obtained from Revelio Labs, to obtain additional (self-reported) personal characteristics of forecasters, such as their educational background and (approximate) age. This merge is feasible because the TGR dataset contains key identifiable information on forecasting analysts, including their full name and employer, enabling a clean and unambiguous match with the Revelio data.

In linking the datasets, we require precise name matches between individuals in the TGR and Revelio datasets, hand-verifying individuals to account for common name abbreviations. To further ensure the accuracy of our matched dataset, we also require that the LinkedIn-provided employment history aligns with the analyst's employer in the equity reports.[11] Additionally, the name-employer match must result in a unique LinkedIn profile within the dataset. That is, we exclude cases where multiple LinkedIn users with the same name and the same employer are linked to the TGR dataset.

4.1.3 Linking Reports to IBES Forecasts

To compare terminal expectations with the same analysts' shorter-horizon forecasts, we complement the TGR dataset with analyst-level forecasts from the IBES unadjusted detail history. We construct

[11] We account for common variations in employer names (e.g., Bank of America vs. BofA) as well as accented characters (e.g., Societe Generale vs. Société Générale).

two shorter-horizon series. The first is a one-year expected EPS growth rate: for each analyst–firm–fiscal-year, we retain one FY1 EPS forecast issued 9 to 14 months before the fiscal year end (selecting the forecast closest to a 12-month horizon when several qualify) and measure expected growth as the forecast relative to the latest realized EPS announced before the forecast date, $(F_1 - A_0)/|A_0|$. The second is the analyst's own long-term growth (LTG) forecast, taken from the IBES LTG detail records. To hold the firm universe fixed across horizons, the decompositions in Section 6.1 restrict both series to the firms covered by the TGR sample. Both growth series are trimmed at the 0.5th and 99.5th percentiles, consistent with the treatment of TGR expectations.

For the matched analysis of Section 6.2, we additionally link IBES forecasts to the individual equity reports. We match on the analyst and the firm, and require the IBES estimate to be filed within five days of the report's publication date; when no estimate is filed in this window, we use the analyst's standing forecast, provided IBES records it as outstanding on the report date.[12] This yields, for a subset of reports, the same analyst's near-term and terminal expectations for the same firm at essentially the same point in time.

4.1.4 Summary Statistics

Table 1 provides summary statistics.

Firms and Coverage (Panel A). Our final sample consists of 41,070 equity reports with extracted information on TGR expectations, from analysts employed by 42 different brokerage houses, covering a total of 7,975 firms across 48 countries. The sample includes expectations data from 5,229 unique forecasting analysts, 2,547 (49%) of whom are successfully matched to LinkedIn. (For forecasters who provide at least five data points in our sample, we take an extra step, including additional manual checking, to link forecasters to LinkedIn, allowing us to match 62% of individual observations in the sample.)

In additional unreported statistics, the average forecasted firm owns assets with a book value of USD 15.8 billion and has an investment rate of 5.55%. These numbers are similar to the IBES universe, where the average firm has assets with a book value of USD 15.2 billion, and an investment rate of 5.3%, suggesting that our sample is comprehensive, representative, and comparable to its commercial counterparts.

Our sample includes two dozen NAICS industry sectors (2-digit), with the eight largest broad sectors accounting for 86% of the total coverage: 35% for manufacturing (NAICS 31-32-33), 19% for information (NAICS 51), 8% for professional services (NAICS 54), 7% for retail trades (NAICS 44-45), 4% for mining and oil & gas (NAICS 21), 4% for transportation (NAICS 48-49), 5% for utilities (NAICS 22), and 4% for finance and insurance (NAICS 52).

[12] An IBES estimate is outstanding between its announcement date and the date the broker last confirmed or revised it (`anndats` to `revdats`), following standard practice (Diether, Malloy, and Scherbina, 2002).

TGR Expectations (Panel B). The average subjective TGR belief in the sample is 2.2%, with the 25th-percentile estimate at 1.5% and the 75th-percentile estimate at 3.0%. By contrast, the average IBES LTG consensus forecast is 12.66%, i.e., an order of magnitude larger.[13] We defer a more in-depth discussion contrasting TGR and LTG expectations to Section 5.1.

Analyst and Country Statistics (Panel C). Regarding individual forecaster characteristics, just above 15% of analysts are female. The majority of analysts are Caucasian (63%) per Revelio's race prediction algorithm. The next largest group is Asian analysts (20%). 27% (21%) of analysts have a Master's (MBA) degree. There is considerable variation in analyst age and birth years, with the 25th-percentile analyst born in 1971 and the 75th-percentile analyst born in 1985.

The average analyst in the sample contributes eight reports and TGR forecasts. There are approximately 110 analysts on average who make forecasts for firms located in a given country, and about 140 analysts per brokerage house.

In unreported statistics, 37% of the reports are produced by analysts based in Europe, 30% in North America (27% from the US), 18% in Asia, 12% in Oceania, 2% in South America, and 1% in Africa. Additionally, 39% of firms have their headquarters located in Europe, 27% in North America (24% in the US), 19% in Asia, 12% in Oceania, 2% in South America, and 1% in Africa (also unreported). As a result of the broad geographic coverage in our sample, about 1 in 3 equity reports and TGR forecasts come from foreign forecasters, i.e., analysts who forecast the long-run growth of a firm headquartered in a different country than where they are located. Finally, there are approximately 860 equity reports and TGR forecasts for firms located in a given country.

Economic Data Series (Panels D and E). Panels D and E present descriptive statistics on various economic series for the countries in which forecasted firms and analysts are located. There is meaningful variation in economic conditions of both firm and analyst countries, both in the full sample (Panel D) and when restricting to the subsample consisting of foreign equity reports (Panel E). For example, the interquartile range (IQR) in firm-country real GDP and inflation in Panel D is 2.2pp and 1.9pp, respectively. Similarly, the IQR in analyst-country GDP and inflation in Panel E is 1.7pp and 1.4pp, respectively. Throughout the analysis, we also draw on country-level macroeconomic series from the World Bank, measures of bilateral proximity between the analyst's and the firm's country (geographic distance, trade shares, social connectedness, language, and music-consumption overlap), and firm-level fundamentals; Appendix Appendix B defines each variable and its source.

[13] We obtain LTG forecasts from Refinitiv (data item code TR.LTGMedian). We trim the IBES LTG variable at the 0.5th and 99.5th percentiles, consistent with the approach used for the TGR expectations.

4.1.5 Familiarity Proximity Between Analysts and Firms

For the roughly one-third of reports in which the analyst and the firm are located in different countries, we measure the geographic, economic, social, and cultural proximity of the two countries. Geographic distance is the distance between the two capitals, and the bilateral trade share is the two countries' trade as a share of the analyst country's total trade in the report year (CEPII BACI). Three further measures capture familiarity that need not decline with distance. Social connectedness is the Facebook Social Connectedness Index (Bailey et al., 2018), the relative probability that a Facebook user in the analyst's country and one in the firm's country are friends, from the May 2026 release of the country-level index.[14] Language similarity is an indicator for the two countries sharing the same main language, defined as the official language, the most spoken official language when there are several unless English is among them, and the most widely spoken language where there is no official language, based on the language profiles compiled from Wikipedia. Cultural affinity is the overlap in music consumption: the share of the two countries' weekly Spotify Top-200 charts that they have in common, averaged over all weeks available for both countries (up to 502 weeks per pair) and scaled from 0 to 100.[15] Appendix Figures B.5, B.6, and B.7 map each measure for four benchmark countries—South Africa, Argentina, Australia, and the United Arab Emirates—and show that connected, same-language, and culturally similar partners include both proximate and distant countries, so that the familiarity splits used in Section 7.2 are not distance splits in disguise. Appendix Appendix B defines each variable and its source.

5 The Relevance of TGR Expectations

This section establishes that TGRs capture meaningful beliefs about very long-run firm growth. We proceed in three steps. First, we show that TGRs differ sharply from IBES LTG, the principal longer-horizon expectation series used in prior research (Section 5.1). Second, we show that the two measures are calibrated to different segments of the growth term structure: LTG is most closely associated with realized growth at its stated three- to five-year horizon, whereas TGR remains strongly associated with cross-sectional growth differences at horizons extending up to thirteen years (Section 5.2). Finally, we show that TGRs are not mechanically contaminated by investors' expected returns (Section 5.3), making them reasonable proxies of long-run growth expectations rather than as a residual that rationalizes observed prices.

[14] The index does not cover mainland China, Hong Kong, or Macau, which are otherwise absent from the social-connectedness split.

[15] Spotify does not operate in mainland China; we proxy China's chart with the Hong Kong chart, consistent with the treatment of Hong Kong-headquartered firms in the main data.

5.1 TGR and IBES LTG Expectations

As noted in Section 4.1.4, TGR and IBES LTG differ sharply in their levels. The gap between the two series averages 10.45 percentage points and persists throughout the sample period (Table 1 and Panel A of Figure 1).[16] The average TGR is 2.2%, close to economy-wide and mature-industry long-run growth rates—levels that a firm could plausibly sustain in perpetuity. The average LTG forecast is 12.66%, far above any growth rate that a firm could sustain forever.

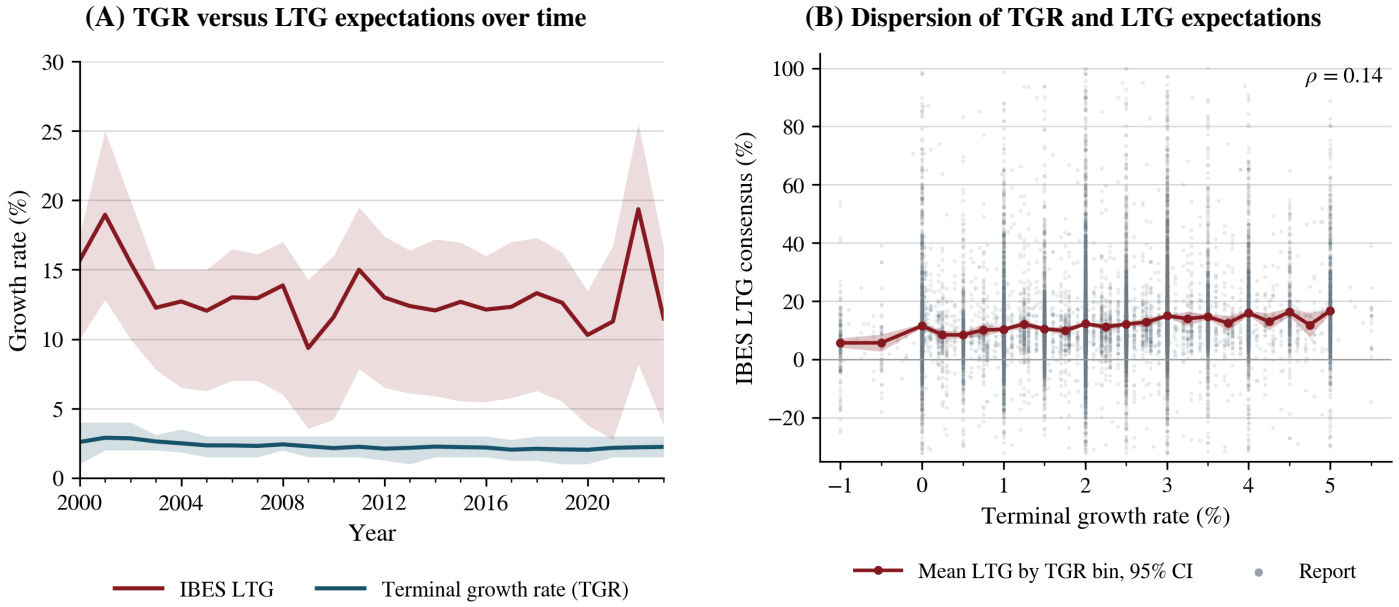


Figure 1: TGR versus IBES LTG expectations This figure compares analysts' terminal growth rate (TGR) expectations with the IBES long-term growth (LTG) consensus for the 25,167 reports in which both are available, corresponding to the "IBES overlap" sample in Table 1. Both series are expressed in percent and trimmed at the 0.5th and 99.5th percentiles. Panel A plots the annual mean and interquartile range of TGR and LTG. Panel B plots each report's TGR against the median LTG consensus for the corresponding firm-year. The red line reports mean LTG within 0.25-percentage-point bins of TGR, centered on the round values at which TGRs cluster; shaded areas report 95% confidence intervals. The statistic ρ is the report-level correlation between the two series. Variable definitions appear in Appendix Appendix B.

The two series also share little cross-sectional variation. The correlation between the TGR reported by an analyst and the IBES LTG consensus for the corresponding firm-year is statistically significant but economically weak, at only 0.14 (Panel B of Table 1). Panel B of Figure 1 illustrates this limited overlap. The binned means slope upward, indicating that the measures contain some common information, but the association is diffuse: as TGR ranges from −1% to 5.5%, mean LTG rises from approximately 6% to 17% against substantial underlying dispersion. Indeed, the standard deviation of LTG is 13.9 percentage points, eleven times the 1.3-percentage-point standard deviation

[16] Relatedly, Appendix Figure B.2 shows that TGRs are systematically below historically realized long-run growth rates in every year of our sample, whereas the LTG expectations examined in prior work are systematically above historical growth rates (Nagel and Xu, 2022; Bordalo et al., 2024).

of TGR.

TGRs are also coarse and slow-moving. They cluster at round values and, in 63% of cases in which the same analyst covers the same firm in consecutive years, remain unchanged. Such persistence is natural for beliefs about growth in perpetuity, since little arriving news is informative about growth in the distant future. It nevertheless raises two empirical questions: whether these slow-moving beliefs remain calibrated to firm-specific differences in subsequently realized long-run growth, and whether their coarseness and infrequent revision mechanically contribute to the forecaster component in our variance decomposition. We address the first question next. We return to the second after presenting the decomposition, using tests designed specifically to isolate the role of coarseness and persistence.

5.2 TGR Expectations and Subsequent Realized Earnings Growth

We evaluate the calibration of TGR and LTG by relating each measure to firms' subsequently realized earnings growth at different horizons. This comparison provides a direct test of whether the two expectation series correspond to different segments of the growth term structure.

Figure 2 first compares the annual mean TGR with average realized earnings growth from $t+3$ to $t+8$.[17] This is our preferred horizon because the modal DCF model ends its explicit forecast period after three years and applies the TGR to all subsequent cash flows. The two series move closely over time: both peak in the early 2000s and then decline steadily. Their levels nevertheless differ, with mean TGR approximately 3.2 percentage points below mean realized growth.

The aggregate comovement in Figure 2 could arise from broad long-run economic trends rather than firm-specific information. Panel A of Table 2 addresses this possibility by examining whether TGRs vary with firm fundamentals after absorbing common country×industry×year conditions. Analysts assign TGRs that are 0.24 to 0.33 percentage points lower to firms with above-median long-horizon default probabilities—a discount of approximately 11 to 14% relative to the sample mean. TGRs also decline with firms' post-IPO age, consistent with the life-cycle patterns examined in Section 7.1. TGRs are thus not generic macroeconomic assumptions appended to valuation models; they incorporate firm-specific information about the sustainability of long-run growth.

Panel B of Table 2 then examines whether these cross-sectional differences are calibrated to subsequently realized earnings growth. We consider three horizons: annualized growth from t to $t+5$, from $t+3$ to $t+8$, and from $t+3$ to $t+13$. The final window measures growth over ten years after the explicit forecast period, trading a smaller sample for a more distant measure of firm growth.[18] Each specification includes both TGR and the firm-year IBES LTG consensus. The

[17] We use $t+3$ to $t+8$ to denote the period from the beginning of year $t+3$ through the beginning of year $t+8$. Realized growth is the geometric average annual growth in EBITDA from the Refinitiv Eikon accounting data; Appendix Appendix B provides details.

[18] As Figure B.3 shows, firms' realized growth rates tend to stabilize relatively quickly as they mature. Growth

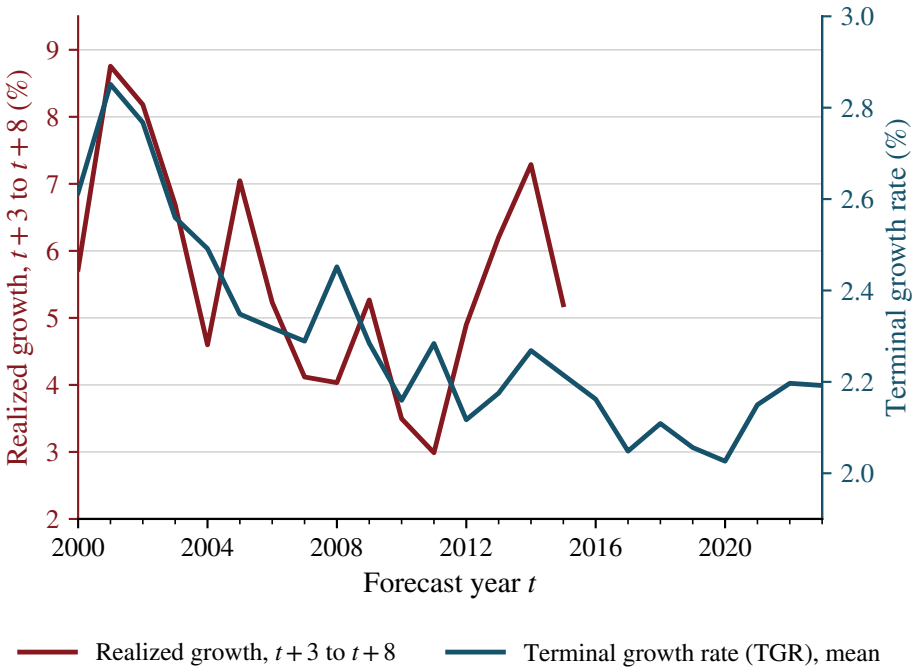


Figure 2: Terminal growth rate expectations and realized growth This figure plots, by forecast year t, the mean terminal growth rate (TGR) expectation across equity reports published in that year against the mean realized earnings growth of the corresponding firms from $t + 3$ to $t + 8$. Realized growth is the geometric average annual growth in EBITDA over the horizon beginning after the modal three-year explicit forecast period. Both series are unweighted averages across reports, so firms covered by more analysts receive more weight. The realized-growth series ends in 2015, the last forecast year for which eight subsequent years of earnings are available. The mean TGR among reports with observed realized growth differs from the full-sample mean by less than 0.1 percentage point in every year from 2004 onward. The full sample covers 2000–2023. Variable definitions appear in Appendix Appendix B.

even-numbered columns additionally include firm-headquarters country×industry×year fixed effects and controls for long-horizon default risk, firm age, and trailing GDP growth in the analyst's country. These specifications compare TGR and LTG while holding fixed broad economic conditions and observable firm characteristics associated with long-run growth.

TGR remains positively associated with realized earnings growth at each horizon. The coefficients without and with fixed effects and controls are, respectively, 0.59 and 0.38 over t to $t + 5$, 0.70 and 0.47 over $t + 3$ to $t + 8$, and 0.49 and 0.30 over $t + 3$ to $t + 13$. Thus, in our preferred controlled specification, a one-percentage-point difference in TGR is associated with a 0.47-percentage-point difference in realized growth after the explicit forecast period. Even over the most distant horizon, the corresponding association remains 0.30 percentage points. We interpret these coefficients as measures of calibration rather than causal effects: firms assigned higher TGRs subsequently realize higher long-run earnings growth on average.

measured up to thirteen years after the forecast date therefore captures an economically important component of long-run firm growth.

The comparison with LTG reveals a sharply different horizon profile. At the stated three- to five-year horizon of IBES LTG (Bordalo et al., 2024), a one-percentage-point difference in LTG is associated with a 0.34- to 0.36-percentage-point difference in realized earnings growth from $t + 1$ to $t + 5$. This association weakens monotonically as the horizon lengthens. It falls to 0.09–0.10 percentage points over $t + 3$ to $t + 8$ and to 0.02–0.06 percentage points over $t + 3$ to $t + 13$. In the controlled specification at the longest horizon, the LTG coefficient is 0.02 and statistically indistinguishable from zero. The corresponding TGR coefficient is 0.30—approximately fifteen times as large—and, once TGR is included, LTG retains virtually no independent association with distant realized growth.

The two measures are thus calibrated to different segments of the growth term structure. LTG is meaningfully associated with realized growth over its stated three- to five-year horizon, but this association fades as the horizon extends. TGR, by contrast, remains calibrated to cross-sectional differences in firm growth well beyond the explicit forecast period, conditional on LTG, observable firm fundamentals, and common country-industry-year conditions.[19] Together with their distinct levels and limited correlation, these results establish that TGR and LTG measure different horizons of expected firm growth.

5.3 Reverse-Engineering TGR Expectations from Prices

Before studying the drivers of TGR expectations in Section 6, we address a remaining measurement concern: analysts may reverse-engineer TGRs from market prices rather than form them independently, a possibility at the center of recent debate (cf. Nagel, 2024). Because we observe both the TGR and the discount rate used within the same valuation model, we can derive a direct test of the mechanical relationship that reverse-engineering implies. We find no such relationship and can rule out economically meaningful contamination of TGRs by expected returns.

Reverse-Engineering. The concern is not simply that analysts use information contained in prices. If analysts accurately infer investors' long-run beliefs, reported TGRs would still reveal meaningful beliefs about long-run firm growth. The concern arises when analysts infer beliefs from current prices using discount rates that differ from those used by investors. The resulting growth estimates would then absorb variation in investors' expected returns. In the extreme case in which analysts use a constant discount rate while investors use a time-varying rate and hold constant growth expectations, high prices reflect low expected returns, but analysts mistakenly infer and report high growth beliefs.

The most intuitive empirical test for this possibility is uninformative. Prior work documents

[19] This separation between medium- and long-run expectations is mirrored in realized growth: the correlation between current and historical growth rates declines steadily as the distance between the two periods increases (Figure B.4).

a negative relation between longer-term expectations and future realized returns (e.g., Bordalo et al., 2024), and we replicate this relation for TGR expectations in Panel A of Table 3.[20] This pattern is consistent with analysts reporting high growth expectations when investors' expected and subsequently realized returns are low. It is, however, equally consistent with analysts and investors using the same valuation model and inputs but systematically overstating or understating firms' growth prospects. Overstated growth beliefs would raise prices today and produce low subsequent returns even in the absence of reverse-engineering. Prior work addresses this ambiguity by showing that analyst expectations predict returns after controlling for proxies of time-varying required returns and that investors and analysts use similar valuation frameworks. Because expected returns are measured imperfectly, however, the extent of reverse-engineering remains debated (Nagel, 2024).[21]

A Direct Test. Because we observe all components of analysts' valuation models, we can test the reverse-engineering hypothesis directly. Let p_t denote the observed market price of the firm. Investor beliefs satisfy:

$$p_t - d_t = E \sum_{j=1}^{N} \rho^j \Delta d_{t+1+j} + \frac{\rho^N}{\rho} E\,TGR - \frac{1}{\rho} r_t, \tag{4}$$

where r_t denotes the discount rate implicit in market prices. Analysts use the same two-part short- and long-run valuation structure when forming price targets but rely on subjective beliefs, denoted by *, and a subjective, time-varying discount rate r_t^a:

$$p_t^a - d_t = E^* \sum_{j=1}^{N} \rho^j \Delta d_{t+1+j} + \frac{\rho^N}{\rho} E^* TGR - \frac{1}{\rho} r_t^a. \tag{5}$$

The first term in each equation captures cash-flow expectations over the explicit forecast horizon, while the second captures expected growth beyond that horizon.[22]

[20] The negative relation between TGR expectations and realized returns obtains both without fixed effects and in specifications including firm and firm HQ country×year fixed effects. It also holds when controlling for analysts' subjective discount rates—the inputs used in DCF models to discount future cash flows—and when using only the components of the subjective cost of equity as predictors of realized returns (Décaire and Graham, 2026).

[21] Bordalo et al. (2024) show that LTG expectations retain their negative relationship with realized returns after controlling for the surplus-consumption and consumption-wealth ratios in Campbell and Cochrane (1999) and Lettau and Ludvigson (2001), respectively, and the SVIX^2 index in Martin (2017). In addition, dividends, earnings, and cash flows predict returns in market data (Golez, 2014; Sabbatucci, 2022; Hillenbrand and McCarthy, 2024; Pruitt, 2024); analysts' subjective expected returns and earnings growth inform changes in their subjective valuations (Décaire and Graham, 2026) and market prices (Chen, Da, and Zhao, 2013; Delao and Myers, 2021); and the CAPM is commonly used to estimate expected equity returns by both investors (Guercio and Tkac, 2002; Barber, Huang, and Odean, 2016; Berk and van Binsbergen, 2016) and analysts (Brav, Lehavy, and Michaely, 2005; Andrei, Cujean, and Wilson, 2023; Décaire and Graham, 2026). Relatedly, Décaire and Graham (2026) show that analysts reporting TGRs use time-varying discount rates and that these estimates track realized future returns on average.

[22] Décaire and Graham (2026) show that analysts typically apply a single discount rate to all projected cash flows within a given DCF model. This rate is not fixed: it changes as analysts revise their valuation models over time.

Under reverse-engineering, analysts start from the observed market price, implying $p_t^a = p_t$, but may solve the pricing identity using a discount rate that differs from the rate implicit in prices, $r_t^a \neq r_t$. Let $\epsilon \equiv r_t^a - r_t$ denote this discount-rate wedge. Because the market price is taken as given, the wedge must be absorbed by analysts' growth beliefs. Allowing fraction ϕ to load onto near-term growth and $1-\phi$ onto terminal growth, the observed price can be written as:

$$p_t - d_t = \underbrace{\left(E \sum_{j=1}^{N} \rho^j \Delta d_{t+1+j} + \phi \cdot \frac{1}{\rho}\epsilon \right)}_{\text{implied short-run beliefs}} + \frac{\rho^N}{\rho} \underbrace{\left(E\,TGR + (1-\phi) \cdot \frac{1}{\rho^N}\epsilon \right)}_{= E^*TGR \text{ (implied long-run beliefs)}} - \underbrace{\frac{1}{\rho}\Big(r_t + \epsilon\Big)}_{= \frac{1}{\rho} r_t^a \text{ (analyst discount rate)}} . \tag{6}$$

If the entire wedge is absorbed by short-term growth expectations ($\phi = 1$), reported TGRs remain unaffected. The problematic case is $\phi \neq 1$, in which the wedge enters TGR mechanically. Equation 6 then yields a direct testable implication: after absorbing macroeconomic conditions that jointly affect discount rates and growth expectations, residual variation in analysts' subjective discount rates should covary with their TGRs. The absence of such comovement provides evidence against contamination of TGRs through reverse-engineering.

Panel B of Table 3 implements this test by relating analysts' subjective TGR expectations to their subjective discount rates:

$$E^*TGR_{i,j,t} = \beta r^a_{i,j,t} + \mu_{i,j,t}. \tag{7}$$

The coefficient β captures the total covariance between subjective TGR expectations and discount rates, scaled by the variance of the latter. This covariance may reflect genuine exposure to common economic conditions or mechanical contamination from the discount-rate wedge, whose contribution is $(1-\phi)\frac{1}{\rho^N}\,\mathrm{var}(\epsilon_{i,j,t})$.[23]

Column (1) includes no fixed effects and yields a small, positive, and statistically significant coefficient, consistent with both mechanical contamination and common macroeconomic drivers such as interest rates. Column (2) adds country×year fixed effects to absorb broad macroeconomic conditions and analyst×firm fixed effects to isolate variation within analyst–firm pairs over time—the level at which reverse-engineering would occur. The estimated coefficient becomes effectively zero and is statistically insignificant at conventional levels (p-value > 0.10).[24] The zero is also

[23] Formally, $\mathrm{cov}((1-\phi)\frac{1}{\rho^N}\epsilon_{i,j,t}, \epsilon_{i,j,t}) = (1-\phi)\frac{1}{\rho^N}\,\mathrm{var}(\epsilon_{i,j,t})$.

[24] The conclusion is unchanged when TGR and r_t^a are expressed in logarithmic form, $\ln(1+TGR)$ and $\ln(1+r_t^a)$, thereby allowing for multiplicative rather than additive errors. Our fixed-effects strategy also avoids relying on potentially noisy estimates of the discount rate implicit in prices. Décaire and Graham (2026) show that extracting expected returns from analyst price targets and current market prices does not recover analysts' subjective discount rates; instead, the resulting measure conflates the subjective cost of equity with analysts' assessments of market mispricing.

precise: evaluated against the mechanical benchmark in equation 6, the estimate implies that only a trivial share of any discount-rate wedge can be flowing into reported TGRs.[25] Columns (3) and (4) replicate the analysis using analysts' subjective cost of equity in place of the discount rate and confirm the result.

Interpreting this precise zero as evidence against contamination requires two conditions. First, the wedge $\epsilon_{i,j,t}$ must vary over time within analyst–firm pairs. A fixed wedge $\epsilon_{i,j}$ would be absorbed by analyst×firm fixed effects, leaving our test with no power against a fixed level distortion. This possibility is narrower than the reverse-engineering concern at issue, however. A fixed wedge cannot generate time-varying contamination of TGRs or account for analysts reporting higher TGRs when investors' expected returns fall. It would also require analysts' stated discount rates to track the unobserved discount rates implicit in market prices one-for-one, up to a constant difference. More generally, reverse-engineering evolving market prices with a misspecified discount rate requires time variation in the wedge; a constant adjustment provides no additional margin with which to match prices as they change. Our test is thus designed to detect the time-varying return contamination that motivates the reverse-engineering concern, while remaining silent about fixed analyst–firm level differences.

Second, the allocation parameter ϕ must lie between zero and one. Within this range, $(1 - \phi)$ is nonnegative, so contamination of TGRs cannot cancel across analysts. Cancellation could arise if some analysts had $\phi < 1$ while others had $\phi > 1$, causing $(1 - \phi)$ to take both positive and negative values. We view such systematic cancellation as implausible. Values of ϕ above one require analysts to allocate more than the entire discount-rate wedge to near-term forecasts and then offset that over-adjustment through an opposite distortion in terminal growth. This procedure introduces larger distortions in both short- and long-run beliefs than are necessary to reconcile the model with the observed price. Under any procedure that favors smaller adjustments to reported inputs, an allocation between zero and one is the natural benchmark. Moreover, as ϕ approaches one, any contamination is concentrated in near-term expectations and leaves TGRs largely unaffected.

Under these conditions, equation 6 predicts a direct, nonnegative, and mechanically strong relationship between analysts' subjective discount rates and reported TGRs whenever expected-return variation contaminates long-run growth beliefs. The evidence in Panel B of Table 3 is not merely

[25] The 95% confidence interval rules out coefficients above 0.016, so a one-percentage-point discount-rate wedge moves reported TGRs by at most 1.6 basis points. Under full absorption of the wedge into terminal growth ($\phi = 0$), equation 6 implies a coefficient of $1/\rho^N \approx 1.1$ for the modal three-year explicit horizon, using log-linearization constants ρ of 0.96–0.97. Longer horizons raise this benchmark and tighten the bound. The ratio of the confidence bound to the benchmark caps the absorbed share at $0.016/1.1 \approx 1.5\%$. This calculation assumes that all residual within-analyst–firm discount-rate variation reflects the wedge. Even if the wedge accounts for only one-quarter of that variation, the cap remains below 6%. Genuine negative comovement between growth beliefs and the wedge could in principle mask some absorption, but its sign is theoretically ambiguous, and the positive estimate in Column (1) indicates that omitted common drivers push the coefficient upward rather than downward.

a failure to reject such a relationship: the precisely estimated zero limits the share of the wedge absorbed by TGRs to approximately 1.5% under the benchmark calculation and to less than 6% even if the wedge accounts for only one-quarter of residual within-analyst–firm discount-rate variation. The evidence rules out economically meaningful return contamination of reported TGRs.

Two interpretations remain consistent with the evidence. Analysts may not engage in widespread reverse-engineering, or they may reverse-engineer prices while allocating most or all of the discount-rate wedge to near-term forecasts. Neither interpretation compromises TGRs as measures of long-run growth beliefs. Nor do our results preclude analysts from learning meaningfully from prices. Our tests complement Bordalo et al. (2024) by using the components of the same analyst's valuation model to test the internal mechanical restriction implied by reverse-engineering, rather than relying on proxies for investors' expected returns. We next study the determinants of TGR expectations in Section 6.

6 The Horizon Structure of Forecast Variation

Section 3 isolates the first dimension along which a common anchor can be tested: the forecast horizon. Under a common anchor, persistent differences among forecasters may generate disagreement at finite horizons, but they should weaken as expectations approach the terminal horizon; if anchors are personal, the opposite pattern should emerge. This section implements that test. We decompose the variation in growth expectations into persistent firm, forecaster, and brokerage-house components at three horizons—one year ahead, the three-to-five-year horizon of IBES LTG, and the terminal horizon—and trace how the source of persistent variation changes as the horizon lengthens (Section 6.1). We then hold the set of forecasters, firms, and dates fixed across horizons in a matched sample (Section 6.2), and show that the forecaster component reflects beliefs rather than reporting mechanics (Section 6.3). The source of persistent variation reverses with the horizon: firm effects account for 63% of the explained variation one year ahead, forecaster effects for 61% at the terminal horizon (Table 4).

6.1 Variance Decomposition Across Horizons

Decomposing TGR Expectations. Table 4 decomposes the variation in growth expectations across brokerage houses, analysts, and forecasted firms using an ANOVA-based sequential variance decomposition. Panel A decomposes TGR expectations. In every panel, Columns (1) to (4) use the full sample, excluding firms, analysts, and brokerage houses with fewer than 2 observations, as the respective fixed effects would fully explain the associated variation in expectations; Columns (5) to (8) further restrict the sample to firms with at least ten observations for robustness. Within each block, the columns require at least 10, 25, and 50 observations per forecaster, to further address

potential concerns about overfitting. Because the shares vary somewhat across these sample cuts, we summarize each panel by the average share across its eight columns and report the range where it matters.

We consistently find that the three sets of fixed effects explain a substantial share—between 61% and 72%—of the variation in TGR expectations. Brokerage house fixed effects account for only a small portion of the explained variation, whereas firm fixed effects and in particular analyst fixed effects drive most of the explained variation. Averaged across the eight columns of Panel A, brokerage house fixed effects contribute 8% of the explained variation, whereas analyst and firm fixed effects contribute 61% and 31%, respectively. The ranking is stable across sample cuts: the analyst share ranges from 42% to 76% across the eight columns and exceeds the firm share in every column but the most restrictive (42% versus 44% in Column (4)). Throughout, analyst fixed effects account for the largest share of the explained variation, suggesting that differences across forecasters play a particularly important role in shaping TGR expectations.

Decomposing Shorter-Term Expectations. Panels B and C apply the same decomposition to the two shorter-horizon expectations of the same analysts, IBES LTG forecasts (three to five years ahead) and IBES one-year EPS growth forecasts, restricted to the firms in the TGR sample so that the firm universe is held fixed across horizons. The results reveal notable differences. At the one-year horizon, firm fixed effects account for 63% of the explained variation on average, analyst fixed effects for 32%, and brokerage house fixed effects for 5% (Panel C). The ranking is thus the reverse of the terminal horizon: analyst fixed effects contribute 30 percentage points more of the explained variation for TGR than for one-year expectations (between 20 and 39 percentage points across the eight column pairs), while firm fixed effects contribute correspondingly less. At the same time, persistent individual heterogeneity remains economically meaningful at short horizons, and, similar to the results for TGR expectations, brokerage house fixed effects explain a negligible portion of the short-run belief variation.

To follow how the shares evolve between these endpoints, we turn to IBES LTG forecasts, whose three-to-five-year horizon lies between them. The estimates fall between them as well (Panel B): analyst fixed effects account for 36% of the explained variation and firm fixed effects for 55%. The persistent drivers of belief variation thus shift monotonically from firms to forecasters as the horizon lengthens. Importantly, the rising forecaster share of explained variation reflects forecaster effects accounting for more of the total variation, not firm effects accounting for less: relative to total variation, the contribution of firm effects itself rises slightly with the horizon.

The pattern we document does not depend on scaling by explained variation. Measured against total variation, the analyst share rises from 7% at the one-year horizon to 12% for LTG and 40% for TGR, whereas the firm share moves comparatively little, from 14% to 21%. Total R^2 rises with the horizon as well, from 0.22 to 0.66: terminal beliefs are far more persistent overall, and most of the

additional persistence is forecaster-specific.

Individual forecaster characteristics matter at every horizon but become increasingly important in terminal beliefs. The result reflects a shift in the source of persistent belief variation, not simply greater dispersion at longer horizons: near-term differences primarily reflect the firms being forecast, whereas terminal differences primarily reflect the forecasters evaluating them. Under a common anchor, heterogeneous learning speeds can generate near-term disagreement, but persistent forecaster differences should fade rather than grow with the horizon (equation 3). The opposite pattern points to differences in the anchor itself or in how forecasters map individual firms onto it.

6.2 Holding Forecasters, Firms, and Dates Fixed: The Matched Sample

The comparisons across the panels of Table 4 hold the firm universe fixed but not the forecasters or forecast dates. The one-year and LTG forecasts come from all IBES analysts covering firms in the TGR sample, whereas the TGRs come from equity research reports. The reversal could therefore reflect differences in sample composition rather than the forecast horizon itself. To address this concern, we link each TGR to the same analyst's one-year-ahead IBES EPS growth forecast for the same firm, either issued within five days of the report or outstanding on the report date (Section 4.1.3). The match reduces the sample to 9,711 reports but allows us to repeat the decomposition holding forecasters, firms, and forecast dates fixed across horizons.

Table 5 reports the result. The shift remains large: analyst fixed effects account for 71% of the explained variation in TGR versus 38% in one-year expectations, and firm fixed effects for 25% versus 61%, with a total R^2 of 0.77 versus 0.42. As such, the shift we document in Table 4 reflects the horizon, not differences in who forecasts, which firms they cover, or when the forecast is performed.

The sequential ANOVA decomposition used above attributes explanatory power shared by brokerage houses, analysts, and firms to whichever set enters first. Although our analysis focuses on how these shares change across horizons rather than on their absolute levels, the ordering could in principle affect that change as well. Because there is little economic basis to guide one ordering, Table 5 reports the order-free Shapley attribution (Shorrocks, 2013), which averages each set's marginal contribution across all possible entry orders. The absolute shares change, but the gradual shift at the core of our analysis remains: analyst effects rise from 34% of the explained variation at the one-year horizon to 48% for terminal beliefs, while firm effects decline from 66% to 50%. Thus, as the horizon lengthens, beliefs become increasingly shaped by who forms them.

6.3 Beliefs Versus Reporting Norms

The coarseness and persistence of TGR expectations documented in Section 5.2 raise a natural concern about the decomposition: a large analyst share could in principle reflect reporting conventions rather than beliefs—round-number defaults, repeated use of the same terminal assumption, or employer templates. We address this possibility in four ways.

Order of Inclusion and Overfitting. One limitation of the approach in Table 4 is that the ANOVA-based estimation is not fully indifferent to variations in the sequential order in which the fixed effects are included. As an alternative approach, we estimate standard OLS regressions in which we gradually include the fixed effects one at a time (Table B.1). Panels A and B replicate Column (1) and Column (8) of Panel A from Table 4, respectively, representing the two extreme cases of that table. Columns (1) to (3) present the results when including each set of fixed effects individually, effectively estimating an *upper bound* on the amount of variation explained by the inclusion of the given fixed effects. Intuitively, the included fixed effects may be correlated with other (omitted) variables that also influence TGR expectations (see the related discussion in Ayyagari, Demirgüç-Kunt, and Maksimovic, 2008). The remaining columns add various sets of fixed effects jointly.

Using the alternative OLS approach, brokerage fixed effects alone explain at most 2% of the total variation in TGR expectations in the full sample in Panel A, while analyst fixed effects alone explain up to 47%, and firm fixed effects alone account for up to 45%. Additionally, brokerage and firm fixed effects together explain up to 47% of the variation, while adding analyst fixed effects increases the R^2 by 18 percentage points, bringing the total R^2 to 65%. To further ensure that these results are not driven by overfitting due to the high degrees of freedom in the fixed effects, we also examine the adjusted R^2, which penalizes model complexity. The results remain consistent: adding analyst fixed effects to brokerage house and firm fixed effects increases the adjusted R^2 by 17 percentage points, from 38% to 55%. The conclusions from the OLS results for the narrower sample in Panel B are similar (brokerage and firm fixed effects 59%, all three 66%; adjusted R^2 from 52% to 60%).

Terminal Horizons. Given the above results that analyst fixed effects account for a large share of the variation in TGR expectations, we also test the possibility that this result merely captures differences in the forecast horizon at which terminal growth begins. For the subsample in which we are able to observe the TGR horizon from analysts' reports, we replicate the R^2 analysis including horizon fixed effects. Table B.2 shows that horizon fixed effects contribute little additional explanatory power (0.62 to 0.63), while analyst fixed effects remain highly informative even conditional on brokerage house and firm fixed effects in this subsample (R^2 of 0.77 with or without horizon fixed effects). These findings indicate that dispersion in long-run growth expectations does not simply reflect differences in TGR horizons, but instead reflects stable analyst-specific heterogeneity.

Employer Practices, Switchers, and Round Numbers. Brokerage house fixed effects explain at most 4% of total variation at any horizon in Table 4, indicating that the forecaster component arises primarily within, rather than across, brokerage houses. Employer practices could nonetheless load onto analyst effects if analysts follow house valuation templates or industry views. Table B.10 addresses this possibility in three ways, using Column (1) of Table 4 as the baseline, where analyst fixed effects account for 69% of the explained variation.

First, we replace brokerage house fixed effects with brokerage house × industry × year fixed effects and enter them first, absorbing any house-wide industry view before analyst and firm effects. This specification is *overly* conservative: it attributes to employers much of the economy- and industry-wide variation that would otherwise load on firm effects. Even so, analyst fixed effects account for 44% of the explained variation, exceeding both employer effects (34%) and firm effects (22%). Second, we restrict the sample to analysts who switch brokerage houses, ensuring that each analyst effect is identified across employers. The analyst share remains 56%. Third, we exclude TGRs that are multiples of half a percentage point, removing differences in analysts' propensity to report round numbers. Among the remaining 3,828 reports, the analyst share is again 56%. Together, these tests show that the forecaster component is not an artifact of employer practices, limited identification across brokerage houses, or round-number reporting.

Persistent Differences Versus Sampling Error. A large analyst share could reflect sampling error in the estimated effects rather than persistent forecaster differences (Kline, Saggio, and Sølvsten, 2020). This concern may be greater at longer horizons: de Silva and Thesmar (2024) document that noise in analyst EPS forecasts rises as the horizon lengthens; if the same pattern extends to terminal forecasts, it would increase the sampling variation in estimated analyst effects. We address both concerns by testing whether the same analyst component carries across independently evaluated firms—the second implication of Section 3. Following Chan and Chen (2026), we randomly divide each analyst's firm relationships into two halves and estimate the analyst effects separately. Because each analyst–firm relationship remains in one half, the estimates share no firms within an analyst. Consequently, firm-specific views should generate covariance between them; under independent forecast and estimation errors, that covariance isolates the component that carries across firms.

The analyst effects estimated on the two halves covary strongly (Table B.9). Across 100 random partitions, the covariance is always positive, with a median of 0.402 squared percentage points, implying a persistent analyst-effect standard deviation of 0.63 percentage points—about half the cross-sectional standard deviation of TGRs. The implied full-sample reliability—the share of estimated analyst-effect variance that carries across firms—is 0.59, which bounds the forecaster share of valuation disagreement in Section 8.[26] Reliability rises from 0.19 for one-year expectations

[26] For comparison, McCaffrey et al. (2009) find that sampling error accounts for 30–60% of the variance in estimated teacher effects; Chan and Chen (2026) report roughly 30% for emergency-department physicians.

to 0.33 for LTG and 0.59 for TGR. Because the analyst share also rises with the horizon, correcting for sampling error strengthens the conclusion that the forecaster component carrying across firms grows with the horizon.

Overall, our interpretation does not require frequent belief revisions; it requires only that stable forecaster-specific differences carry across firms. Having established that such a component exists, we next ask when it is strongest and what gives it content: how it varies over the firm life cycle and with the predictability of distant growth, how it relates to analysts' distance from the firms they cover, and how it shows up in the text of their reports (Section 7).

7 When Forecaster Heterogeneity Is Strongest

Section 3 identifies two further implications of individual anchors: the weight on the forecaster-specific component, $\omega_{i,t}$, should rise where common benchmarks provide less guidance about distant firm growth, and its content should be associated with the forecaster's experience. We test each in turn. First, we examine when the component is strongest using the firm life cycle and a direct measure of how well public information predicts distant growth (Section 7.1). Second, we examine what gives the component content.

7.1 Anchor Strength: The Firm Life Cycle

A natural approach to exploring how TGR expectations relate to firm characteristics is to link cross-firm heterogeneity in TGR expectations to factors associated with the firm life cycle, given that products tend to become more standardized and opportunities for innovation and growth diminish as firms age and industries mature (e.g., Klepper, 1996).

Firm Age. We begin by relating TGR expectations to firm age. As Figure 3 shows, forecasters' TGR expectations decline steadily with firm age. The average TGR is 2.4% for firms aged 10 years or less, about 2% for firms aged 35 to 60, and roughly 0.5% for the oldest firms, reflecting diminishing growth opportunities as firms mature. The magnitudes in the figure reflect the relationship between TGR expectations and firm age cohort *after* controlling for firm and year fixed effects.

Industry, Country, and Firm-Specific Drivers. Motivated by this link between TGR expectations and firm age, we next examine when the explanatory power of firm characteristics is most pronounced and how the importance of firm features evolves over the firm life cycle.

In Columns (1) and (2) of Table 6, we find that a firm's industry and headquarters country explain only a small portion of TGR expectations for young firms. Specifically, these factors account for only 17% of the explained variation for firms aged five years or less. Even 20 to 25 years after a firm's IPO, industry and country account for only about 32%. Column (3) reports the complement: the share of explained variation attributable to analyst effects is 79% for firms younger than 15 years



Figure 3: TGR expectations and firms' life cycle This figure plots TGR expectations over firms' life cycle. The solid blue line corresponds to age-group fixed effects obtained from the regression: $TGR_{i,j,t} = k + I_{\text{Age group}} + I_{\text{Firm}} + I_{\text{Year}} + \epsilon_{i,j,t}$, where we add back the regression constant (k) to better illustrate the terminal growth rate level. The pale blue lines correspond to 95^{th} confidence interval of the $I_{\text{Age group}}$ regression estimate, where we also add back the regression constant (k). Firm age-group captures firm age groups measured in 5-year increments such that, e.g., firm age-group = 10 corresponds to firms between 5 and 10 years old.

and 59% at 20 to 25 years. In Columns (5) and (6), we find that in these young firms, firm fixed effects explain a large share of the variation in TGR expectations that is unaccounted for by analyst, brokerage house, industry, and firm country factors, consistent with the evidence in Section 6.1. Firm fixed effects account for 35% of the residual variation for firms aged five years or less and 28% for firms aged 20 to 25 years, or 11% and 9% of total variation.

These patterns reverse as firms mature. The crossover occurs at age 45: for firms aged 40 to 45, industry and country account for 35% of explained variation and analyst effects for 50%; for firms aged 45 to 50, the shares are 67% and 21%. Beyond age 45, industry and country account for two-thirds of explained variation and analyst effects for 16%, weighting age groups by their number of reports. Accordingly, for mature firms, the incremental explanatory power of firm fixed effects over industry and country declines significantly.[27]

This evidence highlights that the role and nature of firm-specific factors in explaining TGR

[27] While we report magnitudes for very old firms, we caution that the sample size for these cohorts is smaller—fewer than 100 observations—possibly raising concerns about overfitting. That said, the patterns observed for older firms align with the progression seen among younger cohorts, where sample sizes average several thousand observations per age group.

heterogeneity evolve systematically over the firm life cycle. We interpret these findings through the lens of a firm life-cycle framework, in which forecasters face significant uncertainty about whether a new firm will achieve long-term success or face near-term bankruptcy.[28] As forecasters gain information about a firm's long-term viability, expectations for average long-run growth increasingly converge with broader country and industry trends (Pástor and Veronesi, 2003). In the context of our conceptual framework, this is the exercise associated with $\omega_{i,t}$: the weight on the forecaster's own component declines as common benchmarks become more informative. This pattern is also difficult to reconcile with mechanical persistence. A reporting convention fixed within analyst should not fade as the firms she covers mature, yet the analyst share falls from 79% to about 20% over the life cycle.

Dispersion in TGR Expectations Over Firms' Life Cycle. The shares above describe how explained variation is split, not whether disagreement itself shrinks. It does. Appendix Table B.3 shows that the within-industry-year dispersion of TGR expectations across forecasters is highest for young firms and declines as firms mature, by as much as 88% for the most mature firms. The pattern holds with and without industry × year fixed effects and when controlling for firm size, and a Bonferroni-adjusted test of equality of the age coefficients is rejected at the 1% level in every specification. Disagreement is thus greatest when firm-specific anchors are weakest, so the falling analyst share applies to a falling total.

Where a Common Anchor is Hardest to Find. In the life-cycle test, firm age proxies for how difficult it is for forecasters to converge on a common benchmark for distant firm growth. We now measure this intuition more directly. For each firm-year, we predict realized growth from $t+3$ through $t+8$ using a set of common information predictors[29]. We estimate the model leave-one-year-out, ensuring that every growth prediction is out of sample. We then regress the log squared out-of-sample error on the same information; the fitted value, $U_{i,t}$, measures the predicted difficulty of forecasting distant growth for that firm-year. Out-of-sample root mean squared error doubles from the most to the least predictable tercile, from 10 to 20 percentage points. Then, we sort firm-years into terciles of U and repeat the decomposition of Table 6 (Appendix Table B.4). Country and industry factors account for 24% of the explained variation in the most predictable tercile but only 14% in the least predictable tercile.

Because U is estimated from a model that includes country and industry indicators, sorting on U partly sorts on how well country and industry themselves predict growth, which could

[28] Firm survival probability over the first 15 years is 58%; conditional survival then rises monotonically—71% from 15 to 30, 80% from 30 to 45, 85% from 45 to 60, and 87% from 60 to 75. As default risk shrinks, extrapolating a mature firm toward industry or economy-wide growth trends becomes increasingly defensible. Survival is computed from Compustat, 1950–2025, treating delistings for bankruptcy, liquidation, and related reasons as exits and dropping firms acquired or taken private from the window.

[29] The IBES consensus LTG, firm age, size, beta, growth in the firm's country, the firm's own past growth, default risk, and industry and country indicators

mechanically raise their share of TGR variation in the predictable terciles. Panel C ranks U within country $\times$ industry, holding country and industry composition fixed across the sort. Each tercile therefore contains the same countries and industries but differs in firm-level predictability. The estimates are nearly identical: the analyst share rises from 70% to 80%.

7.2 Familiarity: Local Experience and Distant Targets

Section 6 established that forecasters differ persistently in their terminal beliefs. A natural question is what gives these differences content. This section examines how locally experienced economic conditions shape long-run expectations, including those for geographically distant firms. Drawing from our conceptual framework, analysts' own economic experience should receive greater weight when direct signals about the firm's environment are weak, particularly when social or cultural familiarity makes that experience transferable to the target.

Panel A of Table 8 shows how personally and locally experienced economic conditions, and specifically recent GDP growth in the forecaster's country of residence, influence forecasters' long-term firm growth expectations, exploiting the granularity of the data on the geographic locations of both firms and the analysts covering them. The sample is restricted to foreign reports, in which the analyst and the firm are located in different countries, and all specifications include analyst-country $\times$ firm-country fixed effects, which absorb any fixed linkage between the two countries. For instance, if analysts are systematically more likely to cover foreign firms with operations in their home country, these fixed effects ensure that local extrapolation is assessed within such pairs. In Column (2), we find that recent GDP growth in the analyst's country is strongly associated with her TGR expectations for foreign firms: a one-percentage-point higher ten-year average growth rate at home raises the TGR by 0.09 percentage points, about the same as a one-percentage-point higher growth rate in the firm's own country (0.10). The relationship thus holds after controlling for recent GDP growth in the firm's country.[30]

Columns (3) to (6) show that local extrapolation is concentrated in forecasts for firms located in geographically distant countries. Column (3) interacts both growth rates with an indicator for above-median distance between the two countries, and Columns (4) and (5) split the sample at that median. Among close pairs, TGRs load on growth in the firm's country (0.18) and not at all on growth in the analyst's country (0.02). Among far pairs, the pattern reverses: the loading on the firm's country falls to 0.04, while the loading on the analyst's home country rises to 0.14. Column (6) adds firm fixed effects and the bilateral trade share, so that the home-country loading is identified

[30] Interestingly, in Column (1), we find no evidence that local inflation influences long-term expectations when recent inflation experiences are included in the regressions instead of recent local GDP growth. This result suggests that the extrapolation mechanism operates primarily through macroeconomic productivity factors rather than inflation. This conclusion aligns with prior evidence from Décaire and Graham (2026), which documents no significant relationship between TGR beliefs and inflation.

from analysts in different countries covering the same firm; the loading on the firm's country goes to zero and the loading on the analyst's country rises to 0.17. Two shifts in opposite directions cannot be produced by shared shocks across neighboring countries. They are consistent with proximity facilitating the processing of information about the firm's context, and with analysts substituting their own experience for that information as the context becomes distant.

Column (7) repeats the far-pair specification of Column (6) for the analyst's own one-year EPS growth forecast. Recalling the smaller role of forecaster effects in short-term expectations (Table 4), we might expect long-run expectations—being more abstract and less tightly anchored to near-term fundamentals—to be more influenced by variation in forecasters' experiences. Consistent with this, the home-country loading for one-year growth is indistinguishable from zero.

Panel B asks what makes home experience transferable to a distant firm. Among far pairs, we interact both growth rates with three indicators of familiarity between the two countries: social connectedness, measured by the Facebook Social Connectedness Index; a shared main language; and cultural affinity, measured by the overlap in the two countries' weekly Spotify charts (Section 4.1.5). Extrapolation from local experience is strongest when the firm's country is socially or culturally familiar. The home-country loading roughly doubles in connected pairs, from 0.10 to 0.29, and by similar amounts for a shared language and for high music overlap, while the loading on the firm's country is unchanged or lower. The pattern appears separately across all three measures and persists after controlling for bilateral trade. Appendix Figures B.5, B.6, and B.7 show that connected, same-language, and culturally similar partners include both proximate and distant countries, so these splits are not distance splits in disguise. Analysts thus project their own experience onto distant settings, and most strongly when those settings are familiar; unfamiliarity alone is not the mechanism.

Two features of the setting give this result its interpretation. First, analysts rarely relocate: fewer than 6% are observed in more than one country over the sample. The tendency to project home-country experience can thus generate persistent differences across forecasters, which is what Section 6 measures. Second, the result is directional. A stable person-specific tendency toward noisy or default reporting cannot produce a loading on home-country GDP growth that rises with distance and with familiarity, so this evidence separates a subjective anchor from heterogeneous shrinkage toward a common default (Enke and Graeber, 2023).

Overall, the evidence in Table 8 extends prior work on local extrapolation (e.g., Kuchler and Zafar, 2019) by documenting extrapolative effects across countries among professional forecasters and for long-run variables. Extrapolation is concentrated among geographically distant firms and, among those, is strongest where social or cultural familiarity makes home experience transferable. Forecasters rely on their own experience where direct information about the target is weakest and where that experience seems to apply.

7.3 Who the Forecaster Is: Characteristics and Backgrounds

Having established the importance of analyst fixed effects in Section 6.1, we next examine the extent to which analyst observables explain this persistent heterogeneity, leveraging our data on forecaster identities and LinkedIn-matched background information.

Panel A of Table 7 reports the total TGR variation explained by firm, brokerage house, and analyst fixed effects across several sample cuts: excluding firms, analysts, and brokerage houses with fewer than two (or five) observations, whose fixed effects would otherwise fully explain the associated variation; restricting to firms with foreign coverage; and excluding domestic coverage. Panel B decomposes the estimated analyst fixed effects into gender, education, race, birth cohort, and country of residence using a sequential ANOVA, with analyst and employer effects separately identified from analysts who switch brokerage houses (Abowd, Kramarz, and Margolis, 1999). Observable characteristics jointly explain 17% to 33% of the variation in analyst fixed effects, almost entirely through two of them: birth cohort and country of residence at the time of the report. The pattern holds in every sample cut. This share is substantially larger than in prior studies of retail investors (cf. Giglio et al., 2021) and households (cf. Malmendier and Nagel, 2011; Bailey et al., 2019; Ben-David et al., 2018; Kuchler and Zafar, 2019; Das, Kuhnen, and Nagel, 2020; D'Acunto et al., 2023), in which demographics typically account for at most 10% of belief variation.

Three features of our setting likely explain the difference. Most analysts forecast multiple firms over several years, so their fixed effects are estimated more precisely. The sample spans analysts from many countries, which separates analyst-country from firm-country effects that homogeneous samples cannot (e.g., affluent U.S. retail investors at Vanguard in Giglio et al., 2021). And long-run expectations are highly persistent, making them primarily cross-sectional objects (Décaire and Graham, 2026), which amplifies the role of stable individual characteristics relative to the more volatile determinants of short-term expectations.

Appendix Table B.5 confirms the ranking with OLS regressions that add the fixed effects one at a time, for Columns (1) and (6) of Panel B. Birth cohort alone explains 8% to 16% of the variation in analyst effects and country of residence alone 10% to 11%, whereas gender, education, and race each explain 1% or less; adding country of residence to a model with all other characteristics raises the R^2 from 0.08 to 0.17 and from 0.18 to 0.30.

The two dimensions that organize the forecaster component are the two that index experienced macroeconomic history: where an analyst lives and when she was born. This gives part of a_j economic content, consistent with the local-extrapolation results of Section 7.2. Observables leave two-thirds or more of the variation in analyst effects unexplained, so the component is person-specific beyond observed cohort and location; models in which beliefs are shaped by cohort experience are nested in it rather than the whole of it. This subsection describes where the component comes from; its existence rests on the horizon and split-sample results of Section 6.

7.4 Interpretation: Qualitative Discussions in Analysts' Reports

Finally, we use the qualitative discussions in analysts' reports to provide complementary evidence on how forecasters interpret shared information. Analysts covering the same firm at the same time observe much of the same public information, but may differ in the considerations they emphasize when forming long-run expectations. Such differences provide a natural form of differential interpretation in the sense of Kandel and Pearson (1995). We therefore ask whether analysts assigning different TGRs to the same firm also differ systematically in how they allocate attention across the information discussed in their reports.

We measure attention using a 50-topic Latent Dirichlet Allocation (LDA) model, similar to Bybee et al. (2024). The model represents each report as a distribution over topics, allowing us to distinguish between the extensive margin—which topics enter an analyst's discussion—and the intensive margin—how much attention the analyst devotes to each topic. We compare reports written by different analysts covering the same firm in the same year. We measure extensive-margin similarity using Jaccard similarity, the share of topics covered by either analyst that are covered by both, and intensive-margin similarity using Cosine similarity, which captures how similarly the two analysts distribute their attention across topics. The Online Appendix (Section Appendix A) describes the text extraction and preprocessing, model estimation and selection, topic content, and robustness to alternative LDA specifications.

The average report concentrates on a relatively small set of themes: its seven largest topics account for 84% of total content. Analysts covering the same firm at the same time nevertheless differ meaningfully in how they discuss this information. On average, two analysts share 38% of the topics covered by either report, while the Cosine similarity of their topic weights averages 0.81.

Panel A of Table 9 asks whether these differences are associated with disagreement in terminal growth beliefs.[31] Differences in which topics analysts cover are not significantly related to TGR disagreement. Differences in the weights analysts place on those topics, by contrast, are systematically associated with their TGR beliefs: lower Cosine similarity is associated with greater TGR disagreement across specifications. The relationship remains negative when we compare the same analyst pairs across firms while controlling for firm and brokerage-pair fixed effects. Thus, analysts estimating different TGRs for the same firm differ less in whether they discuss a given consideration than in the emphasis they place on shared information.

[31] We measure disagreement in TGR expectations between analyst pairs as $TGR_{A,B,i,t} = (TGR_{A,i,t} - TGR_{B,i,t})^2$, where A and B denote the two analysts in the pair, i indexes the forecasted firm, and t denotes the forecast year. Higher values of $TGR_{A,B,i,t}$ indicate greater differences in analysts' TGR beliefs. Since TGR is a rate, squared differences are comparable across firms without normalization. This contrasts with disagreement measures based on levels (e.g., earnings or price targets), which require normalization to account for differences in magnitude. Consistent with the approaches above, we winsorize the disagreement measure at the 0.5th and 99.5th percentiles.

The content of these differences is informative. We group the LDA topics into five categories—macroeconomics, industry, geography, operations, and valuation—and distinguish the first three, which capture broader drivers of firms' distant prospects, from information about firm-level operations and valuation. Panel B of Table 9 shows that the relationship between topic emphasis and TGR disagreement is concentrated in these broader macroeconomic, industry, and geographic topics. Similarity in their emphasis is strongly negatively associated with TGR disagreement, with coefficients ranging from -0.63 to -0.69. By contrast, the corresponding coefficients for operations and valuation topics range from only 0.05 to 0.10 and are statistically insignificant. The association with broader drivers strengthens when similarity in operations and valuation is included simultaneously. Analysts estimating different TGRs for the same firm therefore place different weights on the same macroeconomic, industry, and geographic information, rather than simply discussing different firm operations or valuation mechanics.

Finally, Panel C connects these differences in interpretation to the geographic evidence in Section 7.2. Analysts located farther from one another exhibit greater divergence in their overall topic weights. The same pattern appears when we isolate geographic topics: analysts who are more geographically distant place systematically different weights on geographic information when discussing the same firm at the same time. These results reinforce the interpretation from the preceding exercises. Forecasters exposed to different environments not only extrapolate differently from their local economic experiences; they also differ in how they emphasize shared information when forming long-run beliefs.

Taken together, the textual evidence provides complementary evidence on the content of persistent forecaster differences. Analysts estimating different TGRs for the same firm at the same time place different weights on shared macroeconomic, industry, and geographic information, and those differences in emphasis grow with the distance between the analysts. The evidence thus corroborates differential interpretation as an additional source of person-specific variation in long-run expectations, alongside the role of locally experienced economic conditions documented above. Together, the exercises in this section trace when the forecaster component is largest and what gives it economic content; we next ask how much these persistent forecaster differences matter for valuations.

8 Application to Asset Pricing

This section explores the consequences of the forecaster-specific anchor. We first decompose the disagreement in valuations among analysts covering the same firm in the same year, and ask how much of it is persistent and forecaster-specific and through which input it enters (Section 8.1). We then use a simple calibration to ask whether terminal-anchor differences of the size we observe are

large enough to matter (Section 8.2).

8.1 The Forecaster Share of Valuation Disagreement

For each report, we reconstruct the value of one dollar of current earnings under the analyst's own inputs. In Column (1) of Table 10, the analyst's own terminal growth rate and discount rate enter, with one-year growth fixed at the sample median and years two to five at the firm-year consensus LTG; in Column (2), the analyst's own one-year growth forecast enters as well, on the 2,649 reports with a matched IBES forecast. Because every other element of the valuation is common, differences in reconstructed value among analysts covering the same firm in the same year reflect only differences in the inputs they chose. Panel A repeats the decomposition of Table 4 on the log of this value: analyst effects account for 46% of the explained variation and firm effects for 50%, with brokerage effects again negligible.

Panel B isolates disagreement. Variation within firm-year is 14% of the total variation in reconstructed values; the remainder is variation across firms, as it should be. Of the within-firm-year disagreement, persistent analyst effects account for 40% in Column (1) and 38% in Column (2). Two analysts valuing the same firm at the same time thus disagree in large part for reasons that have nothing to do with the firm: about 40% of their disagreement is the tendency each analyst carries across every firm she values.[32]

Panel C asks through which input this forecaster component enters. We recompute the valuation with one input at a time fixed at its firm-year mean and record the fall in the persistent analyst component; the shares sum to one hundred by construction. In Column (2), 36% of the forecaster component flows through one-year growth, 16% through terminal growth, and 48% through the discount rate. The terminal-growth share is large relative to the disagreement behind it. Analysts covering the same firm differ by about 1 percentage point in terminal growth, against 11 percentage points in next-year growth; terminal growth carries a sixth of the forecaster component because a one-percentage-point difference moves value by roughly 12%, with three-quarters of reconstructed value lying beyond the explicit forecast horizon. The input most shaped by the forecaster is also the one whose effect compounds in perpetuity.

8.2 A Calibration

The decomposition above measures how much valuation disagreement the forecaster component generates; it does not say whether anchor differences of this size are large enough to matter. We close with a simple calibration to illustrate the asset-pricing implications of our results. Specifically,

[32] This share treats the estimated analyst effects as persistent. Shrinking the analyst component by the split-sample reliability of Section 6.3, 0.59, gives a conservative lower bound of about 24%. The bound is approximate, as the reliability is estimated for TGR effects rather than for the valuation effects themselves.

we build on the evidence above on persistent cross-sectional differences in TGR expectations across forecasters and show that the differences observed empirically can generate substantial and realistic levels of valuation disagreement. The two views of Section 3 provide the benchmark. Under a common anchor, forecasters disagree about value because they converge toward the anchor at different speeds; under individual anchors, because they converge to different places. We calibrate each source to plausible magnitudes—considerable heterogeneity in convergence speeds around a common anchor, against terminal-anchor dispersion of half the size we observe within firm-year—and ask which generates more valuation disagreement. If the second matches the first, the anchor differences documented above are not a small correction to a common-anchor world; they are a source of disagreement of the same order as the learning heterogeneity that world relies on. While a full asset-pricing treatment that models equilibrium pricing in the presence of heterogeneous long-run expectations and traces their implications for returns is beyond our scope, the calibration highlights the quantitative potential of this channel.

Setup. We adapt the constant-gain learning framework of Delao, Han, and Myers (2026, DHM) to a cross-section of forecasters who may differ both in how quickly their firm growth expectations converge and in the terminal anchor toward which they converge. DHM use a representative-agent constant-gain model of gradual learning about firms' mean earnings growth. We adapt their framework to distinguish between two sources of disagreement across forecasters: differences in convergence speeds around a common terminal anchor and differences in the terminal anchors themselves.

As in the conceptual framework of Section 3, we write forecaster j's horizon-s expectation as

$$E_t^j[g_{t+s}] = g_{\infty,j} + \rho_j^{s-1}\left(E_t^j[g_{t+1}] - g_{\infty,j}\right), \qquad 0 < \rho_j < 1, \tag{8}$$

where ρ_j governs the speed at which forecaster j's expectations converge and $g_{\infty,j}$ is the terminal growth anchor. A higher ρ_j implies slower convergence toward the terminal anchor, while a lower ρ_j implies faster convergence. Under a common anchor, $g_{\infty,j} = \bar{g}$, heterogeneous convergence speeds can generate substantial disagreement at intermediate horizons, but disagreement vanishes as the horizon lengthens. Heterogeneous terminal anchors instead sustain disagreement in the limit.

Valuations are then obtained from a two-stage discounted cash flow model with a constant discount rate r:

$$P^j = \sum_{s=1}^{T^*} \frac{CF_0 \prod_{h=1}^{s}\left(1 + E_t^j[g_{t+h}]\right)}{(1+r)^s} + \frac{CF_0 \prod_{h=1}^{T^*}\left(1 + E_t^j[g_{t+h}]\right)}{(1+r)^{T^*}} \frac{1 + g_{\infty,j}}{r - g_{\infty,j}}, \quad r > g_{\infty,j}. \tag{9}$$

Calibration. We simulate $J = 1,000$ forecasters. Baseline parameters are $r = 10\%$, $T^* = 10$, near-term expected growth $\bar{g}_1 = 5\%$ (with a cross-forecaster standard deviation of 0.3 pp), and a

persistence distribution $\rho_j \sim \text{Beta}(\mu_\rho = 0.6, \sigma_\rho = 0.20)$ defined on $(0, 1)$. We normalize initial cash flow to $CF_0 = 1$, which is without loss of generality because prices scale linearly in CF_0. In the common-anchor calibration, all forecasters share a terminal anchor $g_{\infty,j} = 2\%$.

Disagreement arises from (i) heterogeneity in ρ_j (convergence speeds) and (ii) modest dispersion in near-term expectations $E_t^j[g_{t+1}]$ (the 0.3 pp noise around $\bar{g}_1$). We hold discount rates fixed across forecasters, so r does not contribute to disagreement, isolating heterogeneity in growth expectations.

In the alternative calibration, we shut down heterogeneity in convergence speeds (i.e., letting $\sigma'_\rho \to 0$) and instead introduce variation in terminal anchors, $g_{\infty,j} \sim \mathcal{N}(0.02, \sigma^2_\infty)$. We discipline terminal-anchor dispersion at $\sigma_\infty = 0.355$ percentage points, half the median within-firm-year standard deviation of TGRs.

Throughout, we assume that differences in short-run expectations, terminal anchors, and convergence speeds are uncorrelated across forecasters, which isolates the role of each channel.

Results. Figure 4 compares valuation disagreement across the two scenarios using the coefficient of variation, the standard deviation of valuations relative to their mean. Despite the modest dispersion in terminal anchors, heterogeneous anchors generate a coefficient of variation in valuations of 4.0 percent, compared with 3.4 percent from considerable heterogeneity in convergence speeds around a common anchor. Modest differences in where expectations converge thus generate as much valuation disagreement as large differences in how quickly they converge. This result reflects the disproportionate valuation weight of distant cash flows. Differences in convergence speeds affect the path of expected growth but vanish at sufficiently long horizons under a common anchor. Differences in terminal anchors instead affect growth in perpetuity. Thus, even modest person-specific differences in terminal growth beliefs can have material valuation consequences because they apply to cash flows in perpetuity.

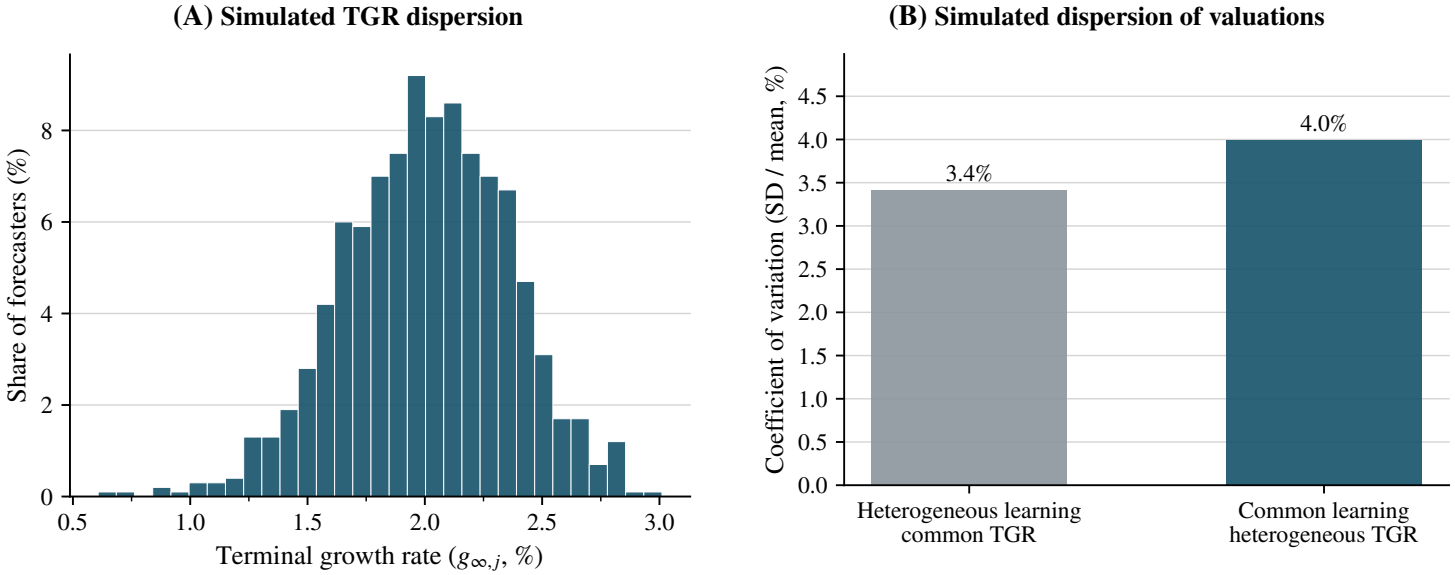


Figure 4: TGR dispersion and disagreement Panel A shows the simulated TGR dispersion in the described calibration with variation in long-run anchors, $g_{\infty,j} \sim \mathcal{N}(0.02, \sigma^2_\infty)$. Panel B plots the coefficient of variation in forecasts for the two calibrations, with heterogeneous learning and common long-run anchors, and with common learning and heterogeneous long-run anchors, respectively.

9 Conclusion

When fundamentals and feedback offer little guidance, do expectations converge on a common anchor or increasingly reflect who forms them? We answer this question using a new dataset of terminal growth rates (TGRs)—professional forecasters' expectations of firm growth in perpetuity—matched to the same forecasters' shorter-horizon forecasts, report text, and backgrounds. Because many forecasters evaluate many of the same firms, we can separate persistent differences across the firms being forecast from persistent differences across the forecasters themselves and trace how their relative importance changes with the horizon.

TGRs capture meaningful beliefs about distant firm growth: they predict realized growth up to thirteen years ahead and are not simply residuals that reconcile valuation models with prices. Yet as the horizon lengthens, the source of persistent variation shifts from firms to forecasters. Forecaster effects are strongest where common benchmarks provide the least guidance: for young firms and for firms whose distant growth is least predictable from public information. These differences are directly related to economic forces. Analysts extrapolate home-country conditions most strongly to geographically distant but culturally familiar markets and weight the same macroeconomic and industry information differently. Because terminal growth assumptions apply in perpetuity, this heterogeneity matters for valuation: forecaster-specific variation accounts for 40% of the valuation disagreement among analysts covering the same firm in the same year.

The results provide the first direct empirical evidence on a primitive over which theories that model or rely on long-run expectations make competing assumptions: whether beliefs converge to a common anchor or retain a persistent person-specific component. Persistent forecaster differences grow rather than fade with the horizon, providing a direct empirical counterpart to models that allow person-specific anchors.

References

Abowd, J. M., F. Kramarz, and D. N. Margolis. 1999. High wage workers and high wage firms. *Econometrica* 67:251–333.

Acemoglu, D., V. Chernozhukov, and M. Yildiz. 2016. Fragility of Asymptotic Agreement under Bayesian Learning. *Theoretical Economics* 11:187–225.

Andrade, P., R. K. Crump, S. Eusepi, and E. Moench. 2016. Fundamental Disagreement. *Journal of Monetary Economics* 83:106–128.

Andre, P., C. Pizzinelli, C. Roth, and J. Wohlfart. 2022. Subjective Models of the Macroeconomy: Evidence From Experts and Representative Samples. *The Review of Economic Studies* 89:2958–2991.

Andre, P., P. Schirmer, and J. Wohlfart. 2026. Mental Models of the Stock Market. *The Quarterly Journal of Economics* .

Andrei, D., J. Cujean, and M. Wilson. 2023. The Lost Capital Asset Pricing Model. *The Review of Economic Studies* 90.

Antoniou, C., C. Frydman, and M. Kilic. 2025. Beyond Firm Characteristics: Forecaster Heterogeneity Explains a Large Portion of Over and Underreaction. *Available at SSRN 5501818* .

Asquith, P., M. B. Mikhail, and A. S. Au. 2005. Information Content of Equity Analyst Reports. *Journal of Financial Economics* 75:245–282.

Ayyagari, M., A. Demirgüç-Kunt, and V. Maksimovic. 2008. How well do institutional theories explain firms' perceptions of property rights? *The Review of Financial Studies* 21:1833–1871.

Bailey, M., R. Cao, T. Kuchler, J. Stroebel, and A. Wong. 2018. Social Connectedness: Measurement, Determinants, and Effects. *Journal of Economic Perspectives* 32:259–280.

Bailey, M., E. Dávila, T. Kuchler, and J. Stroebel. 2019. House price beliefs and mortgage leverage choice. *The Review of Economic Studies* 86:2403–2452.

Bansal, R., and A. Yaron. 2004. Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance* 59:1481–1509.

Barber, B. M., X. Huang, and T. Odean. 2016. Which Factors Matter to Investors? Evidence from Mutual Fund Flows. *The Review of Financial Studies* 29:2600–2642.

Bastianello, F., P. Decaire, and M. Guenzel. 2025a. Mental Models and Financial Forecasts. Tech. rep.

Bastianello, F., P. Decaire, and M. Guenzel. 2025b. Valuation Models. Tech. rep.

Ben-David, I., E. Fermand, C. M. Kuhnen, and G. Li. 2018. Expectations uncertainty and household economic behavior. Tech. rep., National Bureau of Economic Research.

Berk, J., and P. DeMarzo. 2019. *Corporate Finance*. New York.

Berk, J. B., and J. H. van Binsbergen. 2016. Assessing asset pricing models using revealed preference. *Journal of Financial Economics* 119:1–23.

Bertrand, M., and A. Schoar. 2003. Managing with Style: The Effect of Managers on Firm Policies. *The Quarterly Journal of Economics* 118:1169–1208.

Bhamra, H. S., and R. Uppal. 2014. Asset Prices with Heterogeneity in Preferences and Beliefs. *Review of Financial Studies* 27:519–580.

Bharath, S., and T. Shumway. 2008. Forecasting Default with the Merton Distance-to-Default Model. *Review of Financial Studies* 21:1339–1369.

Bordalo, P., N. Gennaioli, R. La Porta, and A. Shleifer. 2025. Finance without Exotic Risk. *Journal of Financial Economics* 173:104145.

Bordalo, P., N. Gennaioli, R. L. Porta, and A. Shleifer. 2024. Belief overreaction and stock market puzzles. *Journal of Political Economy* 132:1450–1484.

Bordalo, P., N. Gennaioli, and A. Shleifer. 2018. Diagnostic Expectations and Credit Cycles. *The Journal of Finance* 73:199–227.

Borovička, J. 2020. Survival and Long-Run Dynamics with Heterogeneous Beliefs under Recursive Preferences. *Journal of Political Economy* 128:206–251.

Brav, A., R. Lehavy, and R. Michaely. 2005. Using Expectations to Test Asset Pricing Models. *Financial Management* 34.

Brealey, R., S. Myers, and F. Allen. 2022. *Principles of Corporate Finance*. New York.

Brock, W. A., and C. H. Hommes. 1997. A Rational Route to Randomness. *Econometrica* 65:1059–1095.

Brunnermeier, M., E. Farhi, R. S. Koijen, A. Krishnamurthy, S. C. Ludvigson, H. Lustig, S. Nagel, and M. Piazzesi. 2021. Perspectives on the future of asset pricing. *The Review of Financial Studies* 34:2126–2160.

Bybee, L., B. Kelly, A. Manela, and D. Xiu. 2024. Business News and Business Cycles. *The Journal of Finance* 79:3105–3147.

Bybee, L., B. Kelly, and Y. Su. 2023. Narrative Asset Pricing: Interpretable Systematic Risk Factors from News Text. *The Review of Financial Studies* 36:4759–4787.

Campbell, J. Y., and J. H. Cochrane. 1999. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy* 107:205–251.

Chan, D. C., and Y. Chen. 2026. The Productivity of Professions: Evidence from the Emergency Department. *American Economic Review* 116:2883–2927.

Chan, D. C., M. Gentzkow, and C. Yu. 2022. Selection with Variation in Diagnostic Skill: Evidence from Radiologists. *The Quarterly Journal of Economics* 137:729–783.

Chaudhry, A. 2025. The Impact of Prices on Analyst Cash Flow Expectations: Reconciling Subjective Beliefs Data with Rational Discount Rate Variation. *Journal of Financial Economics* 171:104095.

Chen, L., Z. Da, and X. Zhao. 2013. What Drives Stock Price Movements? *The Review of Financial Studies* 26:841–876.

Collin-Dufresne, P., M. Johannes, and L. A. Lochstoer. 2016. Parameter Learning in General Equilibrium: The Asset Pricing Implications. *American Economic Review* 106:664–698.

Cookson, J. A., C. Dim, and M. Niessner. 2026. Disagreement on the Horizon. Working paper, University of Colorado Boulder, George Washington University, and Indiana University.

Crump, R. K., S. Eusepi, E. Moench, and B. Preston. 2025. How Do We Learn About the Long Run? Staff Report 1150, Federal Reserve Bank of New York.

D'Acunto, F., D. Hoang, M. Paloviita, and M. Weber. 2023. IQ, expectations, and choice. *The Review of Economic Studies* 90:2292–2325.

Das, S., C. M. Kuhnen, and S. Nagel. 2020. Socioeconomic status and macroeconomic expectations. *The Review of Financial Studies* 33:395–432.

de Silva, T., and D. Thesmar. 2024. Noise in Expectations: Evidence from Analyst Forecasts. *Review of Financial Studies* 37:1494–1537.

Décaire, P. H., and J. R. Graham. 2026. Valuation Fundamentals. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4951338. Working paper.

Decaire, P. H., D. Sosyura, and M. Wittry. 2024. Resolving Estimation Ambiguity. *Working paper* .

Delao, R., X. Han, and S. Myers. 2026. The Cross-section of Subjective Expectations: Understanding Prices and Anomalies. *Working Paper* .

Delao, R., and S. Myers. 2021. Subjective Cash Flow and Discount Rate Expectations. *Journal of Finance* 76:1339–1387.

Delao, R., and S. Myers. 2024. Which Subjective Expectations Explain Asset Prices? *The Review of Financial Studies* 37:1929–1978.

Diether, K. B., C. J. Malloy, and A. Scherbina. 2002. Differences of Opinion and the Cross Section of Stock Returns. *The Journal of Finance* 57:2113–2141.

Ehling, P., A. Graniero, and C. Heyerdahl-Larsen. 2018. Asset Prices and Portfolio Choice with Learning from Experience. *Review of Economic Studies* 85:1752–1780.

Enke, B., and T. Graeber. 2023. Cognitive Uncertainty. *The Quarterly Journal of Economics* 138:2021–2067.

Epstein, L. G., and M. Schneider. 2007. Learning under Ambiguity. *The Review of Economic Studies* 74:1275–1303.

Farmer, L. E., E. Nakamura, and J. Steinsson. 2024. Learning about the Long Run. *Journal of Political Economy* 132:3334–3377.

Gabaix, X. 2020. A Behavioral New Keynesian Model. *American Economic Review* 110:2271–2327.

Giglio, S., M. Maggiori, J. Stroebel, and S. Utkus. 2021. Five Facts about Beliefs and Portfolios. *American Economic Review* 111:1481–1522.

Golez, B. 2014. Expected returns and dividend growth rates implied by derivative markets. *The Review of Financial Studies* 27:790–822.

Graham, J. R., and C. Harvey. 2001. The Theory and Practice of Corporate Finance: Evidence from the Field. *Journal of Financial Economics* 60:187–243.

Guercio, D. D., and P. A. Tkac. 2002. The Determinants of the Flow of Funds of Managed Portfolios: Mutual Funds vs. Pension Funds. *The Journal of Financial and Quantitative Analysis* 37:523–557.

Harrison, J. M., and D. M. Kreps. 1978. Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations. *Quarterly Journal of Economics* 92:323–336.

Hashim, N. A., and N. C. Strong. 2018. Do Analysts' Cash Flow Forecasts Improve Their Target Price Accuracy? *Contemporary Accounting Research* 35:1816–1842.

Hillenbrand, S., and O. McCarthy. 2024. Valuing Stocks With Earnings. *Working Paper* .

Kandel, E., and N. D. Pearson. 1995. Differential Interpretation of Public Signals and Trade in Speculative Markets. *Journal of Political Economy* 103:831–872.

Klepper, S. 1996. Entry, exit, growth, and innovation over the product life cycle. *The American Economic Review* pp. 562–583.

Kline, P., R. Saggio, and M. Sølvsten. 2020. Leave-Out Estimation of Variance Components. *Econometrica* 88:1859–1898.

Kogan, L., S. A. Ross, J. Wang, and M. M. Westerfield. 2006. The Price Impact and Survival of Irrational Traders. *Journal of Finance* 61:195–229.

Kohlhas, A. N., and A. Walther. 2021. Asymmetric Attention. *American Economic Review* 111:2879–2925.

Kuchler, T., and B. Zafar. 2019. Personal experiences and expectations about aggregate outcomes. *The Journal of Finance* 74:2491–2542.

Lettau, M., and S. Ludvigson. 2001. Consumption, aggregate wealth, and expected stock returns. *The Journal of Finance* 56:815–849.

Li, Z., S. Van Nieuwerburgh, and R. Wang. 2026. Understanding Rationality and Disagreement in House Price Expectations. *The Review of Financial Studies* 39:297–342.

Malmendier, U., and S. Nagel. 2011. Depression babies: Do macroeconomic experiences affect risk taking? *The Quarterly Journal of Economics* 126:373–416.

Malmendier, U., and S. Nagel. 2016. Learning from inflation experiences. *The Quarterly Journal of Economics* 131:53–87.

Martin, I. 2017. What is the Expected Return on the Market? *The Quarterly Journal of Economics* 132:367–433.

Martin, I. W. R., and D. Papadimitriou. 2022. Sentiment and Speculation in a Market with Heterogeneous Beliefs. *American Economic Review* 112:2465–2517.

McCaffrey, D. F., T. R. Sass, J. R. Lockwood, and K. Mihaly. 2009. The Intertemporal Variability of Teacher Effect Estimates. *Education Finance and Policy* 4:572–606.

Mukhlynina, L., and K. G. Nyborg. 2020. The Choice of Valuation Techniques in Practice: Education Versus Profession. *Critical Finance Review* 9:201–265.

Nagel, S. 2024. Stock Market Valuation: Explanations, Non-Explanations, and Some Open Questions. Comments Given at the NBER Summer Institute Asset Pricing Group Meeting (Expanded Version).

Nagel, S., and Z. Xu. 2022. Asset Pricing with Fading Memory. *Review of Financial Studies* 35:2190–2245.

Pástor, Ľ., and P. Veronesi. 2003. Stock Valuation and Learning about Profitability. *Journal of Finance* 58:1749–1789.

Patton, A. J., and A. Timmermann. 2010. Why Do Forecasters Disagree? Lessons from the Term Structure of Cross-Sectional Dispersion. *Journal of Monetary Economics* 57:803–820.

Porta, R. L. 1996. Expectations and the Cross-Section of Stock Returns. *Journal of Finance* 51:1715–1742.

Pruitt, S. 2024. Dogs and Cats Living Together: A Defense of Cash-Flow Predictability. *Working Paper* .

Sabbatucci, R. 2022. Are dividends and stock returns predictable? New evidence using M&A cash flows. *Working Paper* .

Shorrocks, A. F. 2013. Decomposition Procedures for Distributional Analysis: A Unified Framework Based on the Shapley Value. *Journal of Economic Inequality* 11:99–126.

Table 1: Summary statistics This table reports summary statistics. The sample includes 40,819 equity reports and terminal growth rate (TGR) expectations for 7,948 firms, obtained from equity reports published between 2000 and 2023. Panel A describes the sample coverage. Panel B reports statistics on the TGR expectations. Panel C reports demographic and other characteristics measured at the individual analyst level (i.e., not at the analyst–year level), along with country-level summary statistics. Panels D and E present various economic data series. Variable definitions appear in Appendix Appendix B.

Panel A: Firms and Coverage	N					
No. Observations (TGR Expectations Sample)	40,819					
No. Firms	7,948					
No. Analysts	5,216					
No. Analysts (LinkedIn-matched)	2,547					
No. Brokerage houses	41					
Panel B: TGR Expectations	Mean	25^{th} Pct.	Median	75^{th} Pct.	Std. Dev.	N
Terminal growth rate (TGR) Expectation $_{i,j,t}$(%)	2.22	1.50	2.00	3.00	1.28	40,819
$TGR_{i,j,t}$(%) (IBES Overlap)	2.21	1.50	2.00	3.00	1.27	25,167
$IBES\ LTG_{i,j,t}$(%)	12.66	5.54	10.30	16.94	13.94	25,167
Correlation(IBES LTG, TGR)	0.14					
Panel C: Analyst and Country Statistics	Mean	25^{th} Pct.	Median	75^{th} Pct.	Std. Dev.	N
Female	0.16	0.00	0.00	0.00	0.36	2,547
Asian	0.19	0.00	0.00	0.00	0.40	2,547
Black	0.03	0.00	0.00	0.00	0.16	2,547
Caucasian	0.63	0.00	1.00	1.00	0.48	2,547
Hispanic	0.04	0.00	0.00	0.00	0.18	2,547
Master degree	0.27	0.00	0.00	1.00	0.44	2,547
MBA	0.21	0.00	0.00	0.00	0.41	2,547
Birth year	1977.51	1971.00	1977.00	1985.00	9.38	1,531
No. Reports by analyst	7.83	1.00	3.00	7.00	16.17	5,217
No. Analysts by country	115.71	8.00	42.00	87.00	253.17	45
No. Analysts by brokerage house	138.05	22.00	42.00	179.00	183.47	42
1(Foreign report)	0.31	0.00	0.00	1.00	0.46	40,519
No. Reports by country	850.38	18.00	245.50	583.00	2111.55	48
Panel D: Economic Data Series—Full Sample	Mean	25^{th} Pct.	Median	75^{th} Pct.	Std. Dev.	N
Firm country						
Real GDP growth $_{i,t}$ (%)	2.54	1.55	2.29	3.74	3.36	37,846
Inflation $_{i,t}$ (%)	2.56	1.26	2.00	3.16	2.50	37,726
10-year Treasury yield $_{i,t}$ (%)	3.15	1.68	2.91	4.10	2.31	39,410
Panel E: Economic Data Series—Foreign Reports	Mean	25^{th} Pct.	Median	75^{th} Pct.	Std. Dev.	N
Analyst country						
Real GDP growth $_{j,t}$ (%)	2.00	1.45	2.29	3.12	3.70	11,894
Inflation $_{j,t}$ (%)	2.47	1.45	2.16	2.85	1.92	11,883
10-year Treasury yield $_{j,t}$ (%)	2.98	1.42	2.85	4.13	2.05	12,417
Firm country						
Real GDP growth $_{i,t}$ (%)	2.10	0.96	2.21	3.82	3.58	11,699
Inflation $_{i,t}$ (%)	2.42	0.90	1.83	3.05	2.75	11,587
10-year Treasury yield $_{i,t}$ (%)	3.02	0.97	2.75	4.11	2.76	11,732

Table 2: Economic relevance of TGR expectations Panel A studies the relationship between TGR expectations and default risk, measured using the Merton Model adjustment from Bharath and Shumway (2008), as well as firm age. The unit of observation is at the firm i, analyst j, and forecast year t level. The sample period is 2000–2023. Panel B regresses ex-post realized firm earnings growth at different horizons on analysts' terminal growth rate expectations and IBES consensus LTG variable. The dependent variable is the geometric average of the annual earnings growth rates experienced by firm i over five- or ten-year periods at various horizons (t to $t+5$, $t+3$ to $t+8$, and $t+3$ to $t+13$, respectively). *Terminal growth rate*$_{i,j,t}$ is the analyst's terminal growth rate expectation, measured in percent. $LTG_{i,t}$ is the IBES consensus LTG belief in percent at time t for firm i. Variable definitions appear in Appendix Appendix B. The regressions are estimated using ordinary least squares, and the standard errors (in parentheses) are heteroskedastic-consistent and clustered at the firm level. Significance levels are shown as follows: * = 10%, ** = 5%, *** = 1%.

Panel A: TGR, Default, and Firm Age	Terminal growth rate (%)$_{i,j,t}$					
	(1)	(2)	(3)	(4)	(5)	
Default risk (10-year horizon)$_{i,t}$	-0.26***			-0.33***		
	(0.03)			(0.03)		
Default risk (20-year horizon)$_{i,t}$		-0.24***			-0.31***	
		(0.03)			(0.03)	
ln(Firm post-IPO age)$_{i,t}$			-0.14***	-0.16***	-0.16***	
			(0.02)	(0.03)	(0.03)	
Firm Country*Industry*Year FE	Yes	Yes	Yes	Yes	Yes	
Observations	20,890	20,890	21,697	13,758	13,758	
F-statistics	91.60	83.27	41.54	55.10	51.99	
R^2	0.36	0.36	0.33	0.37	0.37	
Panel B: Realized Earnings Growth	$(\Pi_{t}^{t+5}(1+g_{i,t})^{1/5}-1)$		$(\Pi_{t+3}^{t+8}(1+g_{i,t})^{1/5}-1)$		$(\Pi_{t+3}^{t+13}(1+g_{i,t})^{1/10}-1)$	
	(1)	(2)	(3)	(4)	(5)	(6)
Terminal growth rate (%)$_{i,j,t}$	0.59***	0.38*	0.70***	0.47*	0.49***	0.30**
	(0.13)	(0.21)	(0.14)	(0.27)	(0.11)	(0.15)
IBES Median Consensus LTG (%)$_{i,t}$	0.36***	0.34***	0.10***	0.09**	0.06***	0.02
	(0.02)	(0.04)	(0.02)	(0.04)	(0.02)	(0.03)
Default risk (20-year horizon)$_{i,t}$		-4.28***		-1.83**		0.10
		(0.73)		(0.78)		(0.66)
ln(Firm post-IPO age)$_{i,t}$		-1.57***		-1.75**		-1.09**
		(0.55)		(0.69)		(0.51)
10-year avg. real GDP growth (Analyst)$_{j,t}$		0.00		-0.25		-0.26*
		(0.17)		(0.19)		(0.13)
Firm Country*Industry*Year FE	No	Yes	No	Yes	No	Yes
Observations	15,242	6,721	12,582	5,438	7,112	2,850
F-statistics	154.09	26.59	22.03	3.26	13.52	2.75
R^2	0.10	0.52	0.01	0.44	0.01	0.56

Table 3: TGR expectations, discount rates, and realized returns Panel A examines the relationship between one-year realized total equity returns, TGR expectations, and analysts' subjective discount-rate estimates. We measure 52-week realized total equity returns using data from Refinitiv Eikon. The subjective equity premium is constructed as $\beta \times ERP$, and the subjective risk-free rate corresponds to the rate analysts use when estimating their cost of equity. Panel B examines the relationship between TGR expectations and subjective discount-rate estimates. The subjective cost of equity is defined as $r_f + \beta \times ERP$, combining the subjective risk-free rate and the subjective equity risk premium. All subjective inputs are drawn directly from analysts' valuation reports. Variable definitions appear in Section Appendix B. All specifications are estimated by ordinary least squares. Standard errors, reported in parentheses, are heteroskedasticity-consistent and clustered at the firm level. Statistical significance is denoted as follows: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Panel A: Realized Returns	**Realized Total Equity Returns (%)$_{i,t+1}$**			
	(1)	(2)	(3)	(4)
TGR (%)$_{i,j,t}$	-0.72**	-1.52***	-0.95*	-2.20***
	(0.31)	(0.33)	(0.55)	(0.81)
Subjective Discount Rate (%)$_{i,j,t}$	1.97***	1.09***		
	(0.23)	(0.32)		
Subjective Risk Premium (%)$_{i,j,t}$			1.67***	1.67***
			(0.40)	(0.60)
Subjective Risk-Free Rate (%)$_{i,j,t}$			1.84***	-1.63**
			(0.55)	(0.77)
Firm FE	No	Yes	No	Yes
Firm Country × Year FE	No	Yes	No	Yes
Observations	24,907	22,852	5,454	4,163
F Statistics	38.17	12.48	7.86	7.03
R^2	0.01	0.47	0.01	0.60
Panel B: TGR and Discount Rate	**TGR$_{i,j,t}$**			
	(1)	(2)	(3)	(4)
Subjective Discount Rate (%)$_{i,j,t}$	0.12***	0.00		
	(0.01)	(0.01)		
Subjective Cost of Equity (%)$_{i,j,t}$			0.08***	-0.02
			(0.01)	(0.01)
Firm Country × Year FE	No	Yes	No	Yes
Firm × Analyst FE	No	Yes	No	Yes
Observations	39,239	23,624	8,399	4,035
F Statistics	449.50	0.01	92.98	2.53
R^2	0.03	0.80	0.02	0.86

Table 4: Variance decomposition of TGR, LTG, and short-run expectations (ANOVA) This table presents a sequential (Type I) ANOVA decomposition of analysts' growth expectations into brokerage-house, analyst, and firm fixed effects, entered in that order. Panel A decomposes TGR expectations, Panel B IBES long-term growth (LTG) forecasts, and Panel C IBES one-year-ahead EPS growth forecasts. In every panel, we exclude firms, analysts, and brokerage houses with fewer than two observations, as their respective fixed effects fully explain the associated variation, and brokerage houses not associated with any analyst who switched employers during the sample period (2000–2023), as the effects of these brokerage houses cannot be separated from those of their analysts (Abowd, Kramarz, and Margolis, 1999). Columns (1) to (4) use the full sample meeting these criteria and columns (5) to (8) further exclude firms with fewer than ten observations; within each group, the columns restrict the sample to analysts with at least the number of reports indicated in the header. For each set of fixed effects, the table reports its share of the explained variation (its sequential sum of squares relative to the model sum of squares), its share of total variation (the same sum of squares relative to the total sum of squares, equal to the total R^2 times the share of explained variation), and its degrees of freedom. Panels B and C use the IBES detail file and include all IBES analysts covering the firms in the TGR sample, with one observation per analyst, firm, and fiscal year; analysts and brokerage houses are identified by their IBES codes. The one-year EPS growth forecast is the analyst's FY1 EPS forecast issued between nine and fourteen months before the fiscal period end (the forecast closest to one year), divided by the absolute value of the most recent realized EPS of the firm announced before the forecast, from the IBES actuals file, minus one. The LTG forecast is the analyst's own IBES long-term growth forecast. Both panels use the same observations, those for which the analyst reports both forecasts, so that the two horizons are compared on identical samples; both forecasts are trimmed at the 0.5th and 99.5th percentiles. Variable definitions appear in Appendix Appendix B.

	Full sample				Obs per firm ≥ 10			
Reports per analyst	All	≥ 10	≥ 25	≥ 50	All	≥ 10	≥ 25	≥ 50
	Panel A: TGR							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Brokerage house FE								
Share of explained R^2 (%)	3	4	6	14	5	6	8	18
Share of total variation (%)	2	3	4	10	3	4	5	12
Degrees of Freedom	39	38	35	27	38	37	34	27
Analyst FE								
Share of explained R^2 (%)	69	63	58	42	76	71	65	46
Share of total variation (%)	46	41	40	30	47	43	41	31
Degrees of Freedom	3,493	1,022	390	102	2,675	950	379	101
Firm FE								
Share of explained R^2 (%)	27	33	36	44	19	23	27	36
Share of total variation (%)	18	22	25	32	12	14	17	24
Degrees of Freedom	5,032	4,502	3,534	2,168	1,160	1,150	1,112	839
Total R^2	0.66	0.66	0.68	0.72	0.62	0.61	0.63	0.67
Observations	36,203	27,124	18,339	8,928	21,580	16,811	11,617	5,541
	Panel B: IBES LTG							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Brokerage house FE								
Share of explained R^2 (%)	10	9	9	7	11	12	12	8
Share of total variation (%)	4	3	3	2	4	4	3	3
Degrees of Freedom	444	353	254	146	438	348	250	144
Analyst FE								
Share of explained R^2 (%)	47	35	27	20	55	42	33	25
Share of total variation (%)	18	12	9	7	19	13	10	8
Degrees of Freedom	7,039	2,526	966	308	6,854	2,523	965	307
Firm FE								
Share of explained R^2 (%)	43	56	64	73	34	46	56	67
Share of total variation (%)	16	19	21	26	12	14	16	21
Degrees of Freedom	3,855	3,624	3,136	2,272	2,352	2,330	2,195	1,748
Total R^2	0.37	0.34	0.33	0.36	0.35	0.30	0.30	0.32
Observations	91,365	72,792	48,896	26,913	84,833	68,177	46,066	25,494
	Panel C: 1-Year Expected EPS Growth Rate							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Brokerage house FE								
Share of explained R^2 (%)	4	5	6	4	5	6	7	5
Share of total variation (%)	1	1	1	1	1	1	1	1
Degrees of Freedom	445	354	253	146	438	347	248	144
Analyst FE								
Share of explained R^2 (%)	47	31	21	15	56	39	26	18
Share of total variation (%)	12	7	5	4	12	7	5	4
Degrees of Freedom	7,041	2,526	966	308	6,857	2,523	965	307
Firm FE								
Share of explained R^2 (%)	48	63	74	81	39	55	67	77
Share of total variation (%)	12	14	16	20	9	10	13	17
Degrees of Freedom	3,864	3,631	3,142	2,279	2,352	2,330	2,194	1,747
Total R^2	0.24	0.21	0.22	0.25	0.22	0.18	0.19	0.22
Observations	91,356	72,843	49,049	27,080	84,709	68,118	46,105	25,573

Table 5: Variance decomposition of TGR and short-run expectations—matched sample This table repeats the variance decomposition of Table 4 on a matched sample in which the TGR and the one-year-ahead EPS growth expectation come from the same analyst, for the same firm, at the same time. The sample comprises the equity reports in our TGR sample for which the analyst's own IBES FY1 EPS forecast for the firm can be matched to the report, issued within five days of the report date or, failing that, the analyst's standing FY1 forecast, provided IBES records it as outstanding on the report date; the one-year growth expectation is this forecast divided by the absolute value of the most recent realized EPS announced before it, minus one, trimmed at the 0.5th and 99.5th percentiles. This removes the composition differences between the two horizons in Table 4, where the short-run forecasts come from all IBES analysts covering the same firms: here the two expectations are decomposed on identical observations and forecasters. We apply the sample filters of Table 4: firms, analysts, and brokerage houses with fewer than two observations are excluded, as are brokerage houses not associated with any analyst who switched employers during the sample period (Abowd, Kramarz, and Margolis, 1999). For each expectation, the *Sequential* column reports the Type I ANOVA decomposition with fixed effects entered in the order brokerage house, analyst, firm, as in Table 4; the *Shapley* column reports the order-free attribution that averages each set's marginal contribution to the R^2 over all six entry orders (Shorrocks, 2013). Both attributions are computed on the same sample, so the total R^2, the number of observations, and the degrees of freedom are identical within each pair of columns. For each set of fixed effects, the table reports its share of the explained variation and its degrees of freedom. Variable definitions appear in Appendix Appendix B.

	TGR		1-Year Expected EPS Growth Rate	
	Sequential (1)	Shapley (2)	Sequential (3)	Shapley (4)
Brokerage house FE				
Share of explained R^2 (%)	4	2	1	1
Share of total variation (%)	3	1	0	0
Range across entry orders (%)		0–3		0–0
Degrees of Freedom	33	33	33	33
Analyst FE				
Share of explained R^2 (%)	71	48	38	34
Share of total variation (%)	55	37	16	14
Range across entry orders (%)		18–57		12–16
Degrees of Freedom	1,292	1,292	1,292	1,292
Firm FE				
Share of explained R^2 (%)	25	50	61	66
Share of total variation (%)	19	38	26	28
Range across entry orders (%)		19–58		25–30
Degrees of Freedom	1,829	1,829	1,829	1,829
Total R^2	0.77	0.77	0.42	0.42
Observations	9,711	9,711	9,711	9,711

Table 6: TGR expectations over firms' life cycle This table analyzes how the factors explaining TGR expectations evolve over the firm life cycle. The unit of observation is at the firm i, analyst j, and forecast year t level; the sample period is 2000–2023. Each row estimates the specifications displayed in the table header separately for firms in the indicated five-year age range, measured from the firm's IPO, after excluding firms, analysts, and brokerage houses with fewer than two reports within the age range. Case 1 regresses TGR expectations on firm-country, industry, brokerage-house, and analyst fixed effects in a single stage. Column (1) reports the R^2 of this regression. Columns (2) and (3) decompose this R^2 with a sequential ANOVA in which the fixed effects enter in the order firm country, industry, brokerage house, analyst: column (2) reports the share of the explained variation attributable to firm country and industry together, and column (3) the share attributable to the analyst fixed effects, which, entering last, measures the analyst component net of country, industry, and brokerage-house factors. Case 2 isolates the firm-specific variation that these factors do not span: Stage 1 recovers the residual $\hat{\epsilon}_{i,j,t}$ from the same four-fixed-effect regression, and Stage 2 regresses this residual on firm fixed effects. Column (5) reports the R^2 of the Stage 2 regression, and column (6) provides a conservative back-of-the-envelope estimate of the share of TGR variation explained by firm factors after accounting for country, industry, brokerage-house, and analyst factors, computed as $(1 - R^2_{Column\ 1}) \times R^2_{Column\ 5}$. Columns (4) and (7) report the number of observations in Case 1 and Case 2; the latter is smaller because singleton groups are dropped in the two-stage estimation. Variable definitions appear in Appendix Appendix B. The regressions are estimated using ordinary least squares.

	Case 1				**Case 2**		
Stage 1:	None				$TGR_{i,j,t} = I_{\text{Firm country}} + I_{\text{Industry}} + I_{\text{Analyst}} + I_{\text{Brokerage}} + \epsilon_{i,j,t}$		
Stage 2:	$TGR_{i,j,t} = I_{\text{Firm country}} + I_{\text{Industry}} + I_{\text{Analyst}} + I_{\text{Brokerage}} + \epsilon_{i,j,t}$				$\hat{\epsilon}_{i,j,t} = I_{\text{Firm}} + \upsilon_{i,j,t}$		
	Stage 2 R^2	**Share of R^2 Driven by Country + Industry**	**Share of R^2 Driven by Analyst**	**Obs.**	**Stage 2 R^2**	**Share of R^2 Driven by Firm Residual Factors**	**Obs.**
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
0 <= Firm age < 5	0.68	17%	79%	2,231	0.35	11%	2,009
5 <= Firm age < 10	0.68	18%	79%	3,569	0.34	11%	3,389
10 <= Firm age < 15	0.64	18%	77%	2,850	0.29	10%	2,702
15 <= Firm age < 20	0.68	26%	66%	2,497	0.30	10%	2,376
20 <= Firm age < 25	0.69	32%	59%	2,005	0.28	9%	1,919
25 <= Firm age < 30	0.77	38%	52%	1,257	0.27	6%	1,202
30 <= Firm age < 35	0.74	37%	51%	632	0.17	5%	606
35 <= Firm age < 40	0.79	31%	54%	474	0.24	5%	463
40 <= Firm age < 45	0.75	35%	50%	181	0.08	2%	170
45 <= Firm age < 50	0.90	67%	21%	134	0.00	0%	133
50 <= Firm age < 55	0.88	58%	22%	107	0.13	2%	105
55 <= Firm age < 60	0.76	81%	13%	76	0.05	1%	74
60 <= Firm age < 65	0.91	64%	23%	84	0.00	0%	82
65 <= Firm age < 70	0.91	80%	0%	44	0.02	0%	43
70 <= Firm age < 75	0.86	88%	1%	51	0.11	2%	50
75 <= Firm age < 80	0.38	45%	9%	31	0.13	8%	30
80 <= Firm age < 85	0.88	59%	6%	34	0.00	0%	33
85 <= Firm age < 90	0.89	29%	29%	33	0.01	0%	32
90 <= Firm age < 95	0.97	76%	20%	15	0.00	0%	15
95 <= Firm age < 100	1.00	100%	0%	12	.	.%	.

Table 7: Analyst fixed effects decomposition This table examines the extent to which analysts' observable characteristics explain the persistent analyst component of TGR expectations. The unit of observation is at the firm i, analyst j, and forecast year t level; the sample period is 2000–2023. The sample comprises analysts with LinkedIn-matched characteristics (gender, race, and birth year). Panel A reports the regression of TGR expectations on firm, brokerage-house, and analyst fixed effects; the analyst fixed effects estimated in this regression are the dependent variable of Panel B. Columns (1) to (3) exclude firms, analysts, and brokerage houses with fewer than two observations, as their respective fixed effects fully explain the associated variation in TGR expectations; columns (4) to (6) raise this requirement to five observations. In all columns we exclude brokerage houses not associated with any analyst who switched employers during the sample period, as the effects of these brokerage houses cannot be separated from those of their analysts (Abowd, Kramarz, and Margolis, 1999), and we require that analysts from at least two countries cover each firm country, so that analyst-country effects are identified within firm country. Columns (1) and (4) use the entire respective subsample, columns (2) and (5) restrict the analysis to firms covered by at least one foreign analyst, and columns (3) and (6) retain only the foreign reports on these firms. Panel B reports, with one observation per analyst, a sequential (Type I) ANOVA decomposition of the estimated analyst fixed effects into five sets of indicators entered in the order shown: gender (equal to 1 if the analyst is female), education (a Master's degree, an MBA, or neither), race (White, Asian, Hispanic, or other), birth cohort (birth year indicators), and analyst country (the country from which the analyst reports most often); each set's share of the explained variation and its degrees of freedom are reported. Variable definitions appear in Appendix Appendix B. Significance levels are shown as follows: * = 10%, ** = 5%, *** = 1%.

Panel A: Fixed effects	**Terminal growth rate$_{i,j,t}$**					
Sample filters	At least 2 observations per firm, brokerage house, and analyst			At least 5 observations per firm, brokerage house, and analyst		
	Full Sample (1)	Has foreign coverage (2)	Exclude domestic coverage (3)	Full Sample (4)	Has foreign coverage (5)	Exclude domestic coverage (6)
Constant	2.17*** (0.01)	2.10*** (0.01)	2.12*** (0.01)	2.11*** (0.01)	2.04*** (0.01)	2.07*** (0.01)
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Brokerage FE	Yes	Yes	Yes	Yes	Yes	Yes
Analyst FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	12,530	5,561	3,587	8,141	4,027	2,655
R^2	0.69	0.70	0.74	0.64	0.67	0.71
Panel B: Anova decomposition	**Analyst FE**					
Sample filters	At least 2 observations per firm, brokerage house, and analyst			At least 5 observations per firm, brokerage house, and analyst		
	Full Sample (1)	Has foreign coverage (2)	Exclude domestic coverage (3)	Full Sample (4)	Has foreign coverage (5)	Exclude domestic coverage (6)
Gender FE						
Share of total R^2 (%)	0	1	0	1	0	0
Degrees of Freedom	1	1	1	1	1	1
Education FE						
Share of total R^2 (%)	1	2	2	0	1	3
Degrees of Freedom	2	2	2	2	2	2
Race FE						
Share of total R^2 (%)	2	2	2	2	1	4
Degrees of Freedom	3	3	3	3	3	3
Birth cohort FE						
Share of total R^2 (%)	45	32	40	37	49	54
Degrees of Freedom	47	43	43	44	42	41
Analyst country FE						
Share of total R^2 (%)	52	63	56	61	48	40
Degrees of Freedom	37	36	28	30	29	23
Total R^2	0.17	0.22	0.29	0.26	0.33	0.30
Observations	955	588	399	556	394	291

Table 8: Analyst–firm distance, local economic conditions, and extrapolation This table examines the effect of analysts' local economic conditions on their long-run expectations for foreign firms. The unit of observation is at the firm i, analyst j, and forecast year t level. The sample period is 2000–2023, and the sample is restricted to non-domestic reports, in which the analyst's country differs from the firm's country. The dependent variables are *Terminal growth rate*$_{i,j,t}$, the analyst's terminal growth rate expectation in percent, trimmed at the 0.5/99.5 percentiles (columns 1–6 of Panel A and Panel B), and the analyst's one-year-ahead expected EPS growth, computed from her IBES FY1 forecast relative to the firm's most recent announced EPS (column 7). The 10-year average real GDP growth rates for the firm's and the analyst's country are obtained from the World Bank and calculated as 10-year trailing averages; the 10-year average inflation rates are constructed analogously. $I_{i,j}^{far}$ is a binary variable equal to 1 if the geographic distance between the analyst's country and the firm's country is above the sample median, and 0 otherwise; distance between countries is measured as the shortest distance between their capitals. Columns 4–5 split the sample at this median, and columns 6–7 and Panel B restrict it to far pairs. *Trade share* is the bilateral trade share between the two countries in year t (CEPII BACI). Panel B interacts both growth rates with three familiarity indicators: *connected*, equal to 1 if the Facebook Social Connectedness Index—the relative probability of a cross-border friendship link between the two countries' populations—is above the sample median; *same language*, equal to 1 if the two countries share the same main language; and *music overlap*, equal to 1 if the countries' average weekly Spotify Top-200 chart overlap is above the sample median. All specifications include analyst country × firm country fixed effects, which absorb the distance and familiarity indicators; columns 6–7 and Panel B add firm fixed effects and the trade share. Sample sizes in Panel B vary with the coverage of each familiarity measure. Variable definitions appear in Appendix Appendix B. The regressions are estimated using ordinary least squares, and the standard errors (in parentheses) are heteroskedastic-consistent and clustered at the analyst country × year level. Significance levels are shown as follows: * = 10%, ** = 5%, *** = 1%.

Panel A: Local extrapolation							
Dependent variable:	TGR$_{i,j,t}$						1-yr EPS growth$_{i,j,t}$
Analyst–firm pairs:	Foreign						
	All			Close	Far		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
10-year avg. inflation (Firm)$_{i,t}$	-0.00						
	(0.00)						
10-year avg. inflation (Analyst)$_{j,t}$	-0.00						
	(0.00)						
10-year avg. real GDP growth (Firm)$_{i,t}$		0.10***	0.18***	0.18***	0.04*	-0.00	-4.32
		(0.02)	(0.03)	(0.03)	(0.02)	(0.03)	(3.92)
10-year avg. real GDP growth (Analyst)$_{j,t}$		0.09***	0.02	0.02	0.14***	0.17***	-5.00
		(0.03)	(0.04)	(0.04)	(0.04)	(0.04)	(8.45)
× Far$_{i,j}$ (Firm)			-0.14***				
			(0.04)				
× Far$_{i,j}$ (Analyst)			0.11**				
			(0.05)				
Analyst country × firm country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	No	No	No	No	No	Yes	Yes
Trade share control	No	No	No	No	No	Yes	Yes
Observations	11,477	11,575	11,575	5,890	5,685	4,678	1,183
F Statistics	0.24	27.62	18.96	28.40	16.77	12.04	1.05
R^2	0.14	0.15	0.15	0.10	0.19	0.53	0.37
Panel B: Drivers of local extrapolation, far pairs							
Dependent variable:	TGR$_{i,j,t}$						
Familiarity indicator $I_{i,j}^{fam}$ = 1 if:	Facebook SCI ≥ median	Same language	Music overlap ≥ median				
	(1)	(2)	(3)				
10-year avg. real GDP growth (Firm)$_{i,t}$	0.00	0.06	0.05				
	(0.04)	(0.04)	(0.04)				
$\times I_{i,j}^{fam}$	-0.04	-0.14**	-0.10*				
	(0.05)	(0.06)	(0.06)				
10-year avg. real GDP growth (Analyst)$_{j,t}$	0.10**	0.12***	0.10**				
	(0.05)	(0.04)	(0.05)				
$\times I_{i,j}^{fam}$	0.19**	0.25**	0.17**				
	(0.08)	(0.10)	(0.08)				
Bilateral trade share$_{i,j,t}$	0.11**	0.11**	0.11**				
	(0.05)	(0.05)	(0.05)				
Analyst country × firm country FE	Yes	Yes	Yes				
Firm FE	Yes	Yes	Yes				
Familiarity indicator $I_{i,j}^{fam}$	Absorbed by the pair FE						
Observations	4,365	4,678	4,591				
F Statistics	9.08	11.48	8.23				
R^2	0.53	0.53	0.53				

Table 9: Variation in TGR expectations and differences in qualitative discussion of topics This table examines the relationship between differences in textual discussions in equity reports and disagreement in TGR expectations between pairs of analysts. The unit of observation is at the analyst a, analyst b, firm i, and forecast year t level. The sample period is 2000–2023. The dependent variable in Panels A and B, $Disagreement_{a,b,i,t}$, is defined as the squared difference between analyst a's and analyst b's TGR beliefs: $(TGR_{a,i,t} - TGR_{b,i,t})^2$. The key independent variables are the Jaccard and Cosine scores, which measure the similarity levels between the two analysts' textual discussions along the extensive and intensive margins, respectively. Panel C relates the Cosine similarity scores to analyst distance. Variable definitions appear in Appendix B. The regressions are estimated using ordinary least squares, and the standard errors (in parentheses) are heteroskedastic-consistent and clustered at the analyst-pair level. Significance levels are shown as follows: * = 10%, ** = 5%, *** = 1%.

Panel A: Differences in Qualitative Discussion	Disagreement$_{a,b,i,t}$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Jaccard $_{a,b,i,t}$	-0.17	-0.17	-0.11				-0.09
	(0.25)	(0.25)	(0.25)				(0.21)
Cosine $_{a,b,i,t}$				-0.53**	-0.47**	-0.42*	-0.31*
				(0.22)	(0.22)	(0.22)	(0.18)
Year FE	No	Yes	No	No	Yes	No	No
Industry FE	No	Yes	No	No	Yes	No	No
Industry*Year FE	No	No	Yes	No	No	Yes	Yes
Analyst pair FE	No	No	No	No	No	No	Yes
Firm FE	No	No	No	No	No	No	Yes
Brokerage pair FE	No	No	No	No	No	No	Yes
Observations	23,466	23,433	23,400	23,466	23,433	23,400	19,753
F Statistics	0.49	0.50	0.21	5.69	4.60	3.77	1.54
R^2	0.00	0.02	0.07	0.00	0.02	0.07	0.74

Panel B: Differences by Content Type	Disagreement$_{a,b,i,t}$					
	(1)	(2)	(3)	(4)	(5)	(6)
Cosine $^{\text{Distant Drivers}}_{a,b,i,t}$	-0.63**	-0.56**	-0.62**	-0.69***	-0.62**	-0.67***
	(0.25)	(0.25)	(0.25)	(0.25)	(0.25)	(0.25)
Cosine $^{\text{Operations \& Valuation}}_{a,b,i,t}$				0.05	0.09	0.10
				(0.18)	(0.17)	(0.17)
Year FE	No	Yes	No	No	Yes	No
Industry FE	No	Yes	No	No	Yes	No
Industry*Year FE	No	No	Yes	No	No	Yes
Observations	22,927	22,895	22,861	22,426	22,394	22,360
F Statistics	6.49	5.16	6.22	3.90	3.22	3.78
R^2	0.00	0.02	0.07	0.00	0.02	0.07

Panel C: Distance Between Analysts	Cosine$_{a,b,i,t}$			Cosine $^{\text{Geography}}_{a,b,i,t}$		
	(1)	(2)	(3)	(4)	(5)	(6)
$I^{far}_{a,b,t}$	-0.01**	-0.01***	-0.01***	-0.01***	-0.01***	-0.01***
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Year FE	No	Yes	No	No	Yes	No
Industry FE	No	Yes	No	No	Yes	No
Industry*Year FE	No	No	Yes	No	No	Yes
Observations	57,125	57,079	57,053	29,561	29,527	29,493
F Statistics	6.42	9.53	8.31	32.97	29.96	24.87
R^2	0.00	0.01	0.04	0.00	0.01	0.05

Table 10: The forecaster share of valuation disagreement and its sources This table decomposes the disagreement in valuations among analysts covering the same firm in the same year into persistent forecaster-specific differences and apportions that share across the valuation inputs. The valuation is the synthetic value of one dollar of current earnings under the analyst's own inputs. In column (1), the analyst's own terminal growth rate and discount rate enter, with one-year growth fixed at the sample median and years two to five at the firm-year consensus LTG; in column (2), the analyst's own one-year growth forecast enters as well, on reports with a matched IBES forecast. Panel A reports the pooled sequential ANOVA of Table 4, Column (1), with brokerage, analyst, and firm fixed effects entered in that order, and the share of explained variation attributable to each. Panel B reports the sequential ANOVA with firm-year, brokerage, and analyst fixed effects entered in that order, so that the analyst share is the share of within-firm-year variation—disagreement among the analysts covering the same firm in the same year—attributable to persistent analyst effects after brokerage effects. Panel C recomputes the valuation with one input at a time fixed at its firm-year mean and reports the fall in the analyst sum of squares of Panel B, averaged over the two removal orders, as a share of the baseline analyst sum of squares; the shares sum to one hundred because fixing every input at its firm-year mean removes all within-firm-year variation. Samples require the terminal growth rate to lie at least 0.5 percentage points below the discount rate, for each analyst and for the firm-year means of both inputs, one-year growth above −99%, and at least two analysts per firm-year, and apply the sample rule of Table 6.

	Own TGR and r (1)	Own g^1, TGR, and r (2)
Panel A: TGR variance decomposition		
Reports	16,927	2,649
Variance of ln value	0.204	0.368
R^2	0.588	0.675
Brokerage share of explained variation (%)	3.5	3.7
Analyst share of explained variation (%)	46.4	46.4
Firm share of explained variation (%)	50.1	50.0
Panel B: Within-firm-year disagreement decomposition		
Within-firm-year variance of ln value	0.029	0.051
Within-firm-year variance, share of total (%)	14.1	14.0
Persistent analyst effects, share of disagreement (%)	40.3	38.5
Panel C: Share of the persistent forecaster component (%)		
One-year growth	—	36.3
Terminal growth rate	31.5	16.1
Discount rate	68.5	47.5

Online Appendix

The Formation of Very Long-Run Firm Growth Expectations

Paul H. Décaire Marius Guenzel
ASU Wharton

Not For Publication

Appendix A Construction and Validation of Topic Measures

This appendix describes the construction of the topic measures summarized in Section 7.4. We first extract and preprocess the qualitative content of analysts' reports, estimate a Latent Dirichlet Allocation (LDA) model, and then construct measures of similarity in topic coverage and attention across analysts.

Appendix A.0.1 Text Extraction and Preprocessing

For the topic analysis, we extract the full text of each equity report rather than only the quantitative valuation inputs described in Section 4.1.1. We exclude text contained in figures and tables, as well as the legal boilerplate disclaimers appearing at the end of reports.

We use Meta's Llama-2-70b model to address residual extraction problems.[33] We use this refinement only to repair extraction artifacts before applying the remaining text transformations.

We retain reports containing at least 100 words. We then remove verbs and stopwords, lemmatize the remaining words, and organize terms into unigrams and bigrams using their co-occurrence across the corpus. Words that frequently appear consecutively are treated as a single bigram; for example, "supply" and "chain" may enter the vocabulary as "supply chain." Finally, we retain only unigrams and bigrams appearing in at least 0.1% of reports.

Appendix A.0.2 LDA Model

We use LDA to represent each report as a distribution over latent topics, following an estimation strategy similar in spirit to Bybee et al. (2024). Let d index reports, n index tokens within a report, $k \in \{1, \ldots, K\}$ index topics, and V denote the vocabulary size. The model assumes

$$\boldsymbol{\theta}_d \sim \text{Dirichlet}(\alpha \mathbf{1}_K)\,, \tag{A.1}$$

$$\boldsymbol{\phi}_k \sim \text{Dirichlet}(\eta \mathbf{1}_V)\,, \tag{A.2}$$

$$z_{dn} \mid \boldsymbol{\theta}_d \sim \text{Categorical}(\boldsymbol{\theta}_d)\,, \tag{A.3}$$

$$w_{dn} \mid z_{dn}, \{\boldsymbol{\phi}_k\}_{k=1}^{K} \sim \text{Categorical}\big(\boldsymbol{\phi}_{z_{dn}}\big)\,, \tag{A.4}$$

where $\boldsymbol{\theta}_d$ is report d's distribution over topics and $\boldsymbol{\phi}_k$ is topic k's distribution over terms. The estimated topic share $\widehat{\theta}_{dk}$ measures the fraction of report d attributed to topic k. The resulting representation captures both which topics a report discusses and the relative attention assigned to each.

We set $\alpha = 0.05$ and $\eta = 0.02$. These values optimize perplexity in our implementation and are consistent with the standard practice of initializing symmetric Dirichlet hyperparameters on the order of $1/K$. Our results are insensitive to perturbations around these values.

Appendix A.0.3 Selecting the Number of Topics

We estimate models with the number of topics, K, ranging from 10 to 200 in increments of 10. We select among these specifications using three criteria: topic coherence, in-sample perplexity, and manual inspection of the highest-probability terms within each topic. Coherence measures the semantic similarity of the terms assigned to a topic, while perplexity assesses model fit. Manual

[33] For example, PDF extraction sometimes removes the spaces between adjacent words. The language model can infer and restore the intended word boundaries.

inspection allows us to determine whether the resulting topics are economically interpretable and whether additional topics identify new themes or merely divide existing ones.

Figure B.8 reports coherence and perplexity across values of K. Most improvements in both measures occur by 50 topics, after which the gains level off. Manual inspection produces the same conclusion. The 50-topic model yields distinct and largely nonoverlapping economic themes. Increasing K beyond 50 primarily fragments existing themes without capturing substantial new economic content. Appendix Table B.6 reports the highest-probability terms associated with each topic in the preferred specification.

We therefore use the 50-topic model as our baseline and repeat the main analysis using an 80-topic specification. The thematic granularity of the baseline model is also comparable to alternative approaches that define economically interpretable topics ex ante and then group them into broader categories (Bastianello, Decaire, and Guenzel, 2025a). Although the methods differ, both approaches capture analysts' qualitative reasoning at a similar level of abstraction. Table B.8 shows that the section's main results are robust to using 80 topics.

Appendix A.0.4 Topic Coverage and Attention

The estimated model assigns every report a vector of topic weights,

$$\widehat{\boldsymbol{\theta}}_d = \left(\widehat{\theta}_{d1}, \ldots, \widehat{\theta}_{dK}\right), \qquad \sum_{k=1}^{K} \widehat{\theta}_{dk} = 1.$$

The average report concentrates its attention on a limited number of topics. Its seven largest topics account for 84% of total text content, on average. Topics outside this group receive approximately 0.4 percentage points of report attention each.

We classify topic k as *covered* by report d when $\widehat{\theta}_{dk} \geq 0.01$. Among covered topics, analysts allocate an average of 15% of report content to each topic, with an interquartile range of 10% to 20%.

We compare topic allocation across pairs of analysts covering the same firm at the same time. The comparison distinguishes two dimensions of attention: whether analysts discuss the same topics and how much weight they place on the topics they discuss.

Extensive Margin: Topic Coverage. Let A and B denote the sets of topics covered by two reports. We measure overlap using Jaccard similarity:

$$\text{Jaccard}(A, B) = \frac{|A \cap B|}{|A \cup B|}. \tag{A.5}$$

The measure equals zero when the reports cover no common topics and one when they cover identical sets of topics.

The average Jaccard similarity among analyst pairs covering the same firm at the same time is 0.38, with an interquartile range of 0.29 to 0.47. Thus, only 38% of topics covered by at least one member of the average analyst pair appear in both reports. For example, a Jaccard similarity of 0.40 arises if one analyst covers topics 1, 2, 3, and 4, while the other covers topics 1, 2, and 5.

These results indicate that analysts agree on a core set of topics but differ substantially in the broader set of considerations included in their reports.

Intensive Margin: Topic Weights. We measure similarity in topic emphasis using Cosine similarity. Let $\mathbf{w}_A$ and $\mathbf{w}_B$ denote the topic-weight vectors used in the comparison. Cosine similarity is

$$\text{Cosine}(\mathbf{w}_A, \mathbf{w}_B) = \frac{\mathbf{w}_A' \mathbf{w}_B}{\|\mathbf{w}_A\| \, \|\mathbf{w}_B\|}. \tag{A.6}$$

Because topic weights are nonnegative, Cosine similarity ranges from zero to one. A value of one indicates proportional topic weights, while lower values indicate greater differences in relative emphasis.

The average Cosine similarity across analysts covering the same firm at the same time is 0.81, with an interquartile range of 0.73 to 0.93. The corresponding average Cosine dissimilarity is therefore 0.19.

The cross-industry comparison should nevertheless be interpreted cautiously because structural topics—such as valuation methodology, management guidance, and consensus forecasts—appear in most equity reports and mechanically raise similarity across all report pairs. We therefore additionally calculate Cosine similarity using only substantive topics related to macroeconomic conditions, geography, industry conditions, and firm operations. We also report a centered-Cosine specification that removes the corpus-wide average topic distribution before comparing reports.

Appendix A.0.5 Robustness and Interpretation

The similarity patterns are not driven by differences in report length. Both measures are constructed from topics defined by their share of report content rather than raw word counts. Restricting the sample to report pairs whose lengths differ by less than 1% yields an average Jaccard similarity of 0.39 and an average Cosine similarity of 0.83, close to the corresponding baseline values of 0.38 and 0.81.

The results are also robust to estimating the topic model with 80 rather than 50 topics, as reported in Table B.8. Taken together, these exercises show that the topic model captures cross-analyst variation along both the extensive margin—which topics analysts discuss—and the intensive margin—how much weight they place on those topics. The main text uses these measures to examine whether differences in topic attention have systematic economic content and whether they vary with TGR disagreement and the distance between analysts.

Appendix B Variable Definitions

Table A.1: Variable Definitions

Subscript t indicates the forecast year, i a firm, j an analyst, and a, b a pair of analysts covering the same firm in the same year. Rates are in percent. TGR, IBES LTG, one-year expected EPS growth, and Disagreement are trimmed at the 0.5th and 99.5th percentiles.

Variable	Definition
10-year avg. inflation rate $(\text{Analyst})_{j,t}$, $(\text{Firm})_{i,t}$	10-year trailing average of CPI inflation in the analyst's country or the firm's headquarters country (World Bank).
10-year avg. real GDP growth $(\text{Analyst})_{j,t}$, $(\text{Firm})_{i,t}$	10-year trailing average of real GDP growth in the analyst's country or the firm's headquarters country (World Bank).
1-year expected EPS $\text{growth}_{i,j,t}$	The analyst's IBES one-year-ahead EPS forecast relative to the latest IBES actual, $(F_1 - A_0)/\lvert A_0 \rvert$. In Table 5, the forecast is the same analyst's, matched to the equity report.
Analyst $\text{country}_{j,t}$	Country in which the analyst is located when the report is published.
Analyst FE_j	Analyst fixed effect from a regression of TGR on firm, brokerage house, and analyst fixed effects.
Bilateral $\text{trade share}_{i,j,t}$	Trade between the analyst's and the firm's country as a share of the analyst country's total trade (CEPII BACI).
Birth cohort_j	The analyst's year of birth (LinkedIn, via Revelio Labs).
$\text{Cosine}_{a,b,i,t}$	Cosine similarity of the two reports' LDA topic weights over the topics covered by both (50-topic model). The Distant Drivers, Operations & Valuation, and Geography versions restrict to the topics in those categories (Table B.6); the 80 version uses an 80-topic model.
Default risk (10-year, 20-year horizon)$_{i,t}$	Equal to 1 if the firm's probability of default, from the naive Merton model of Bharath and Shumway (2008) at a 10- or 20-year horizon, is above the sample median.
$\text{Disagreement}_{a,b,i,t}$	$(TGR_{a,i,t} - TGR_{b,i,t})^2$.
Education_j	MBA, master's degree, or neither (LinkedIn).
Facebook $\text{SCI}_{i,j}$	Facebook Social Connectedness Index between the analyst's and the firm's country.
Female_j	Equal to 1 if the analyst's predicted gender is female (Revelio Labs).
Firm $\text{age}_{i,t}$	Years since the firm's IPO. Age groups are five-year brackets labeled by their upper bound.
Firm $\text{size}_{i,t}$	Natural logarithm of total assets.
$1(\text{Foreign report})_{i,j,t}$	Equal to 1 if the analyst's country differs from the firm's headquarters country.
$I^{far}_{i,j}$	Equal to 1 if the distance between the analyst's and the firm's country capitals is above the median among foreign reports.
$I^{far}_{a,b,t}$	Equal to 1 if the distance between the two analysts' country capitals is above the median among pairs of analysts in different countries (0 for same-country pairs).
$I^{fam}_{i,j}$	Familiarity indicators: Facebook SCI at or above the median; same dominant language; music overlap at or above the median.
$\text{Jaccard}_{a,b,i,t}$	Jaccard similarity of the sets of LDA topics covered by the two reports, $\lvert A \cap B \rvert / \lvert A \cup B \rvert$, a topic being covered if it accounts for at least 1% of the report (50-topic model; the 80 version uses an 80-topic model).
LTG $(\text{IBES})_{i,t}$	IBES consensus (median) long-term growth forecast at the beginning of the calendar year. $\text{LTG}_{i,j,t}$ is the analyst's own IBES long-term growth forecast.
Music $\text{overlap}_{i,j}$	Overlap between the weekly Spotify Top-200 charts of the analyst's and the firm's country, averaged across weeks, scaled 0–100.
Race_j	The analyst's predicted race (Revelio Labs): White, Asian, Hispanic, or other.

Variable	Definition
Realized earnings growth$_i$	Geometric average annual earnings growth over t to $t+5$, $t+3$ to $t+8$, or $t+3$ to $t+13$.
Realized total equity return$_{i,t+1}$	Total stock return over the 52 weeks following the report (Refinitiv Eikon).
Subjective discount rate$_{i,j,t}$	Discount rate (WACC) of the analyst's DCF model, from the equity report.
Subjective risk-free rate$_{i,j,t}$, risk premium$_{i,j,t}$, cost of equity$_{i,j,t}$	The risk-free rate r_f, the premium $\beta \times ERP$, and their sum $r_f + \beta \times ERP$, from the analyst's DCF inputs in the equity report.
Terminal growth rate (TGR)$_{i,j,t}$	Perpetual growth rate applied beyond the explicit forecast horizon in the analyst's DCF model, from the equity report.
TGR industry dispersion$_{i,j,t}$	Squared deviation of the analyst's TGR from the age group–industry–year mean, scaled by the absolute industry–year mean: $(TGR_{i,j,t} - \overline{TGR}_{k,t}^{\text{Age Group}})^2 / \lvert\overline{TGR}_{k,t}\rvert$.

Appendix C Supplementary Figures and Tables

Appendix Figures

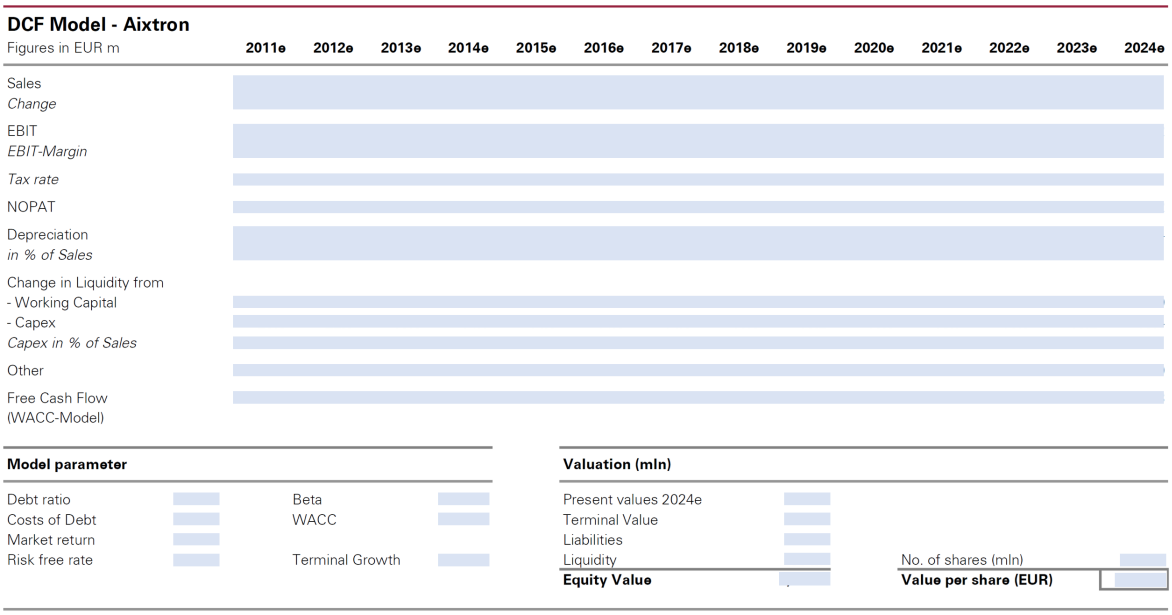


Figure B.1: Example of complete DCF equity report This figure shows a representative example of a discounted cash flow (DCF) valuation model table. This figure is taken from an equity report for Aixtron (Ticker = AIXGn), published by Warburg Research on October 27, 2011. For copyright reasons, all information in the table is redacted.

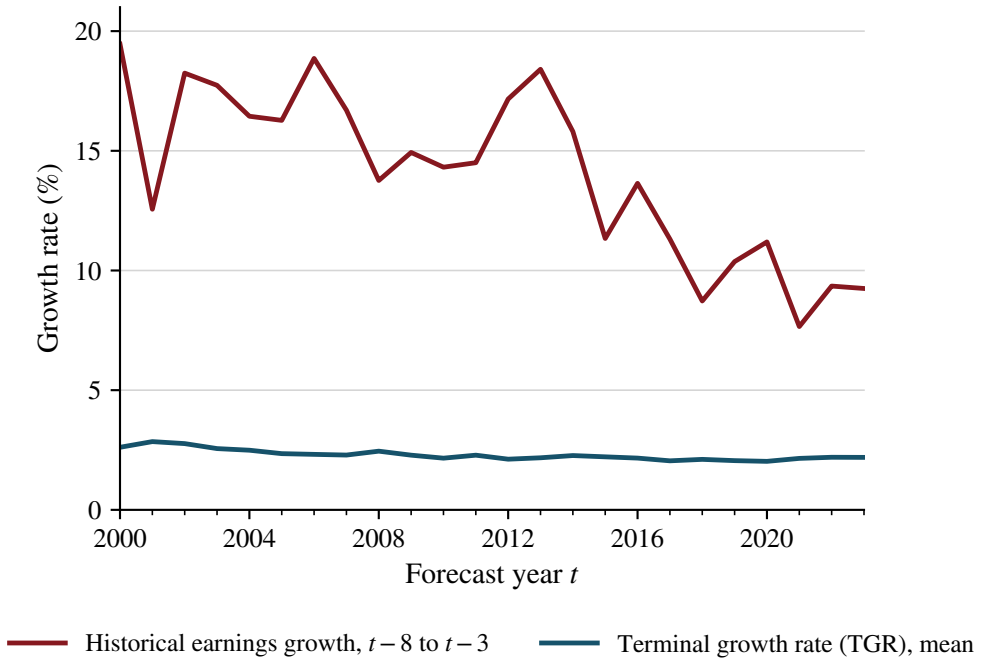


Figure B.2: TGR expectations versus historical returns This figure compares the time trends of historically realized earnings growth at the $t-8$ to $t-3$ horizon with analysts' terminal growth rate expectations at time t. Realized earnings growth corresponds to the geometric average realized between $t-8$ and $t-3$, averaged by year (t). Terminal growth rate is the average terminal growth rate expectation by year. The x-axis denotes years. The y-axis corresponds to rates, in percent (%).

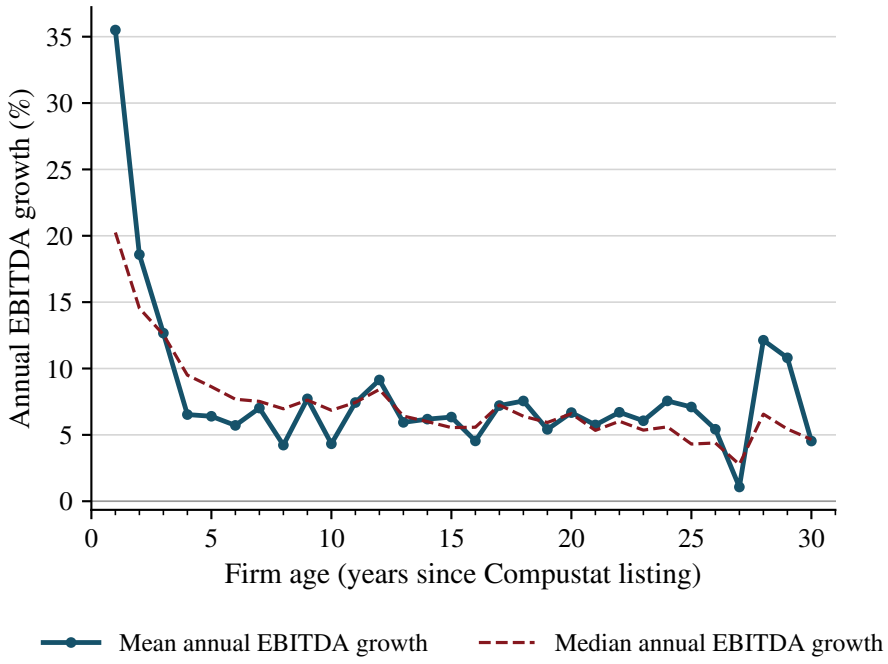


Figure B.3: Annual earnings growth rate and firm age This figure plots firms' realized annual earnings growth over the firm life cycle, motivating the horizons over which we measure realized long-run growth: because growth stabilizes within roughly a decade of listing, realized growth measured up to 13 years out captures a significant component of firms' long-run growth. The sample includes all Compustat firms with available operating earnings before depreciation (OIBDP) that enter Compustat in 1975 or later, so that the firm's first year of coverage marks its listing rather than the start of the database. Firm age is measured in years since the firm's first Compustat year, and annual earnings growth is the year-over-year percentage change in OIBDP, defined for firm-years with positive prior-year earnings and trimmed at the 1st and 99th percentiles; the solid line is its average across firms at each age from 1 to 30, and the dashed line the median. The number of firms per age declines from 15,595 at age one to 979 at age thirty.

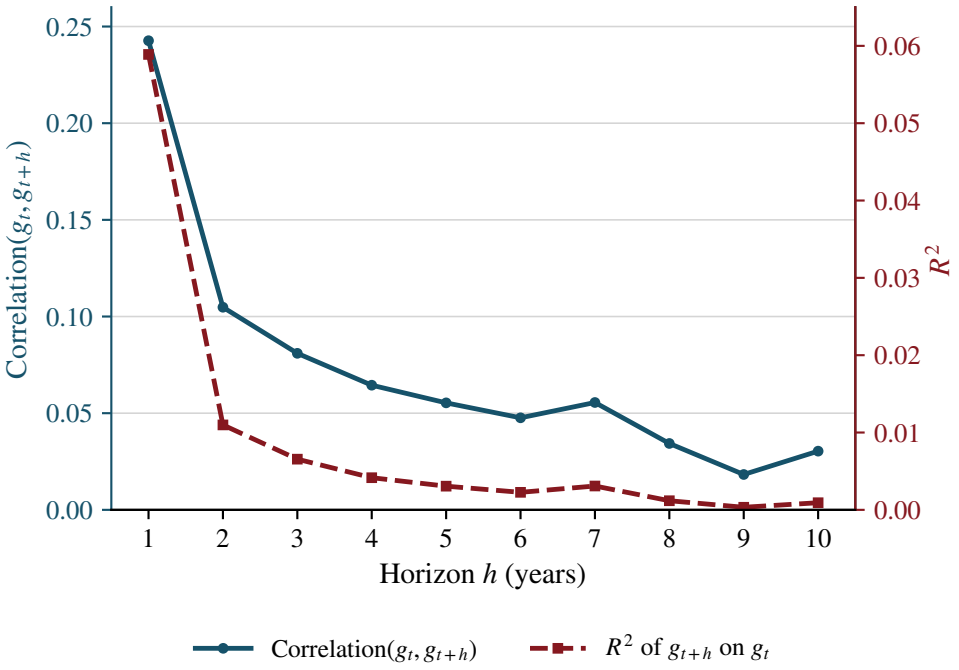


Figure B.4: Correlation between realized growth rates at different horizons This figure plots the average correlation between the realized earnings growth rate at time t and realized earnings growth rate at time $t+h$. Note both t and $t+h$ refer to realized growth rates observed after $t+h$. The blue lines shows the estimated correlation coefficients for different values of h. The red line shows the R^2 when regressing earnings growth at time $t+h$ on earnings growth at time t for different values of h.

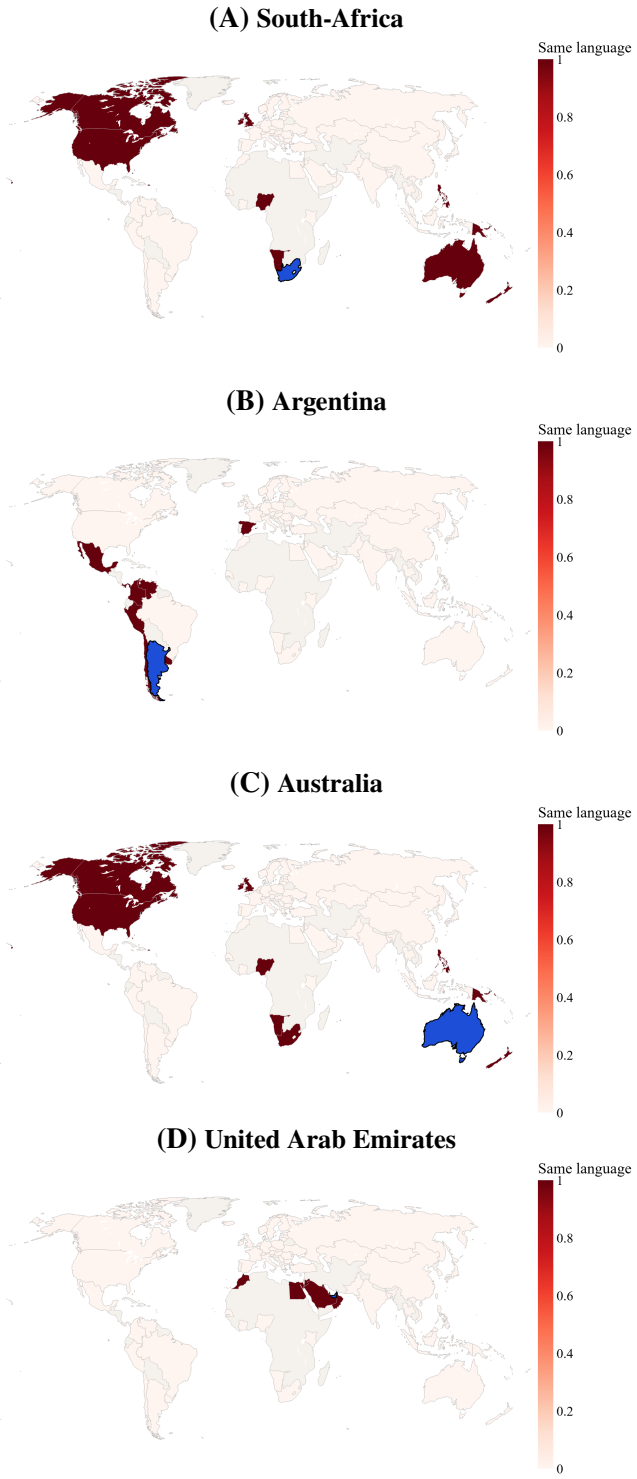


Figure B.5: Language similarity between countries This figure validates the language-similarity split used in Panel B of Table 8 by showing, for four benchmark focal countries, (i) South Africa, (ii) Argentina, (iii) Australia, and (iv) the United Arab Emirates, which of the other countries covered by our sample share their language. The focal country is shown in blue, countries sharing its language in dark red, and covered countries with a different language in light shading; for every focal country, same-language partners include both geographically proximate and distant countries, so the language split is not merely a distance split. We define a country's language as its official language; for countries with multiple official languages, we select the most spoken one, unless English is among them, in which case we prioritize English, and for countries without an official language (e.g., the United States prior to March 1, 2025), we use the most widely spoken language. The similarity measure equals 1 if the focal country and the partner country share the same defined language and 0 otherwise.

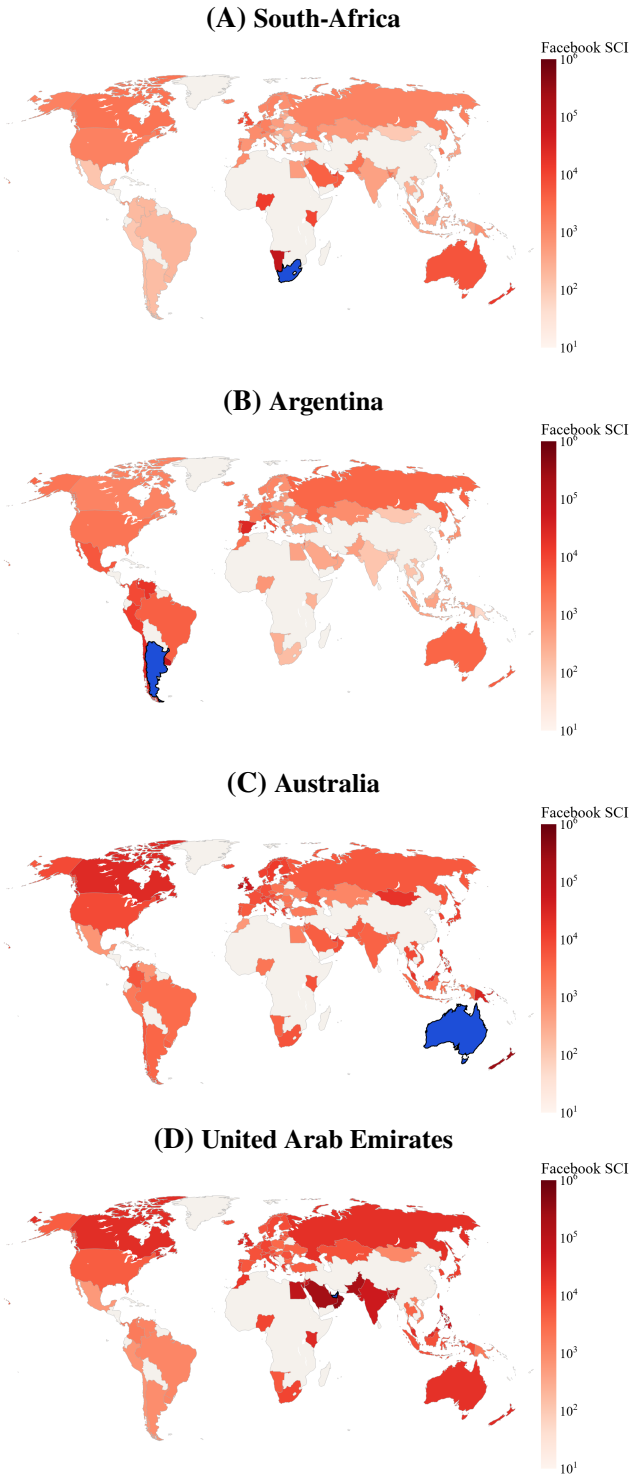


Figure B.6: Facebook social connectedness between countries This figure validates the social-connectedness split used in Panel B of Table 8 by showing, for four benchmark focal countries, (i) South Africa, (ii) Argentina, (iii) Australia, and (iv) the United Arab Emirates, the strength of their Facebook friendship ties with the other countries covered by our sample. The measure is the Facebook Social Connectedness Index (Bailey et al., 2018), the relative probability that a Facebook user in the focal country and a user in the partner country are friends. The focal country is shown in blue, and partner countries are shaded by the index on a logarithmic scale common to the four panels, from 10^1 (lightest) to 10^6 (darkest); countries in grey are not covered by the index (in particular mainland China, Hong Kong, and Macau) or by our sample. For every focal country, strongly connected partners include both geographically proximate and distant countries, so the connectedness split is not merely a distance split. In Table 8, the familiarity indicator equals 1 when the pair's index is at or above the median among foreign analyst–firm pairs.

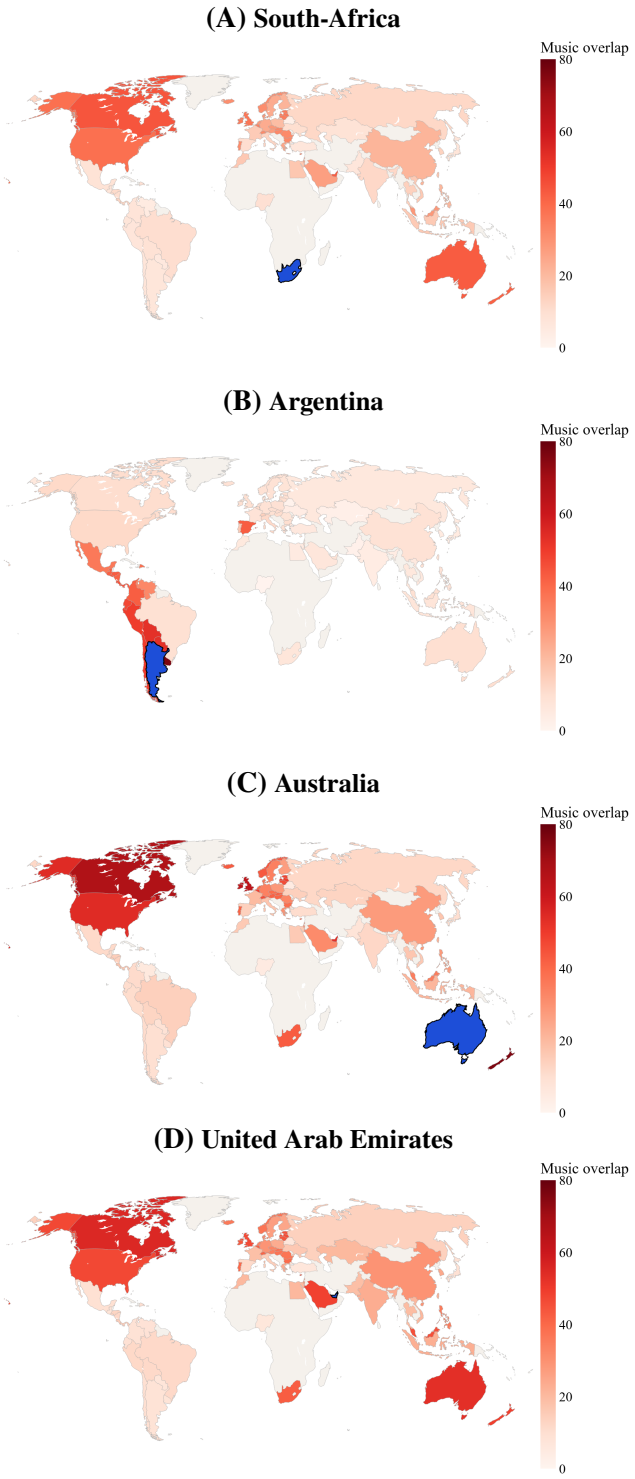


Figure B.7: Cultural affinity between countries This figure validates the cultural-affinity split used in Panel B of Table 8 by showing, for four benchmark focal countries, (i) South Africa, (ii) Argentina, (iii) Australia, and (iv) the United Arab Emirates, the overlap in music consumption between the focal country and the other countries covered by our sample. The measure is the share of the two countries' weekly Spotify Top-200 charts that they have in common, averaged across all available chart weeks, and scaled from 0 (no overlap) to 100 (identical charts). The overlap is computed over up to 502 chart weeks per country pair (median 502, interquartile range 399–502); Spotify does not operate in mainland China, so China's chart is proxied by the Spotify Hong Kong chart, consistent with our treatment of Hong Kong–headquartered firms. The focal country is shown in blue, and partner countries are shaded on a linear scale common to the four panels, from 0 (lightest) to 80 (darkest); countries in grey are not covered by the Spotify charts or by our sample. For every focal country, culturally close partners include both geographically proximate and distant countries – for Argentina, for example, Uruguay but also Spain – so the affinity split is not merely a distance split. In Table 8, the familiarity indicator equals 1 when the pair's overlap is at or above the median among foreign analyst–firm pairs.

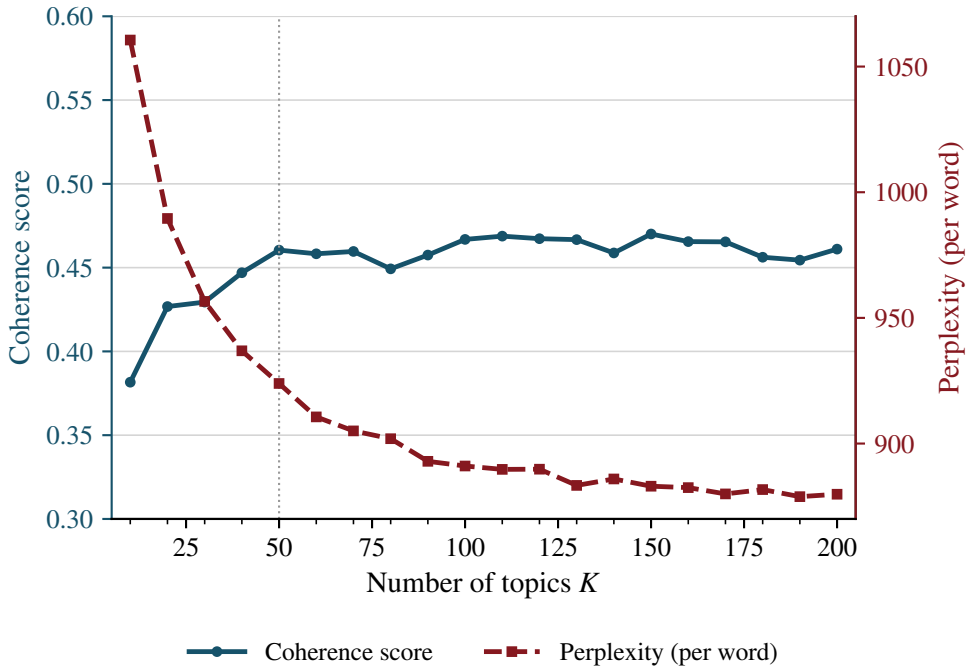


Figure B.8: Coherence and perplexity scores This figure shows the model-selection evidence behind the 50-topic LDA specification of Section 7.4: for models estimated with the number of topics K ranging from 10 to 200 in increments of 10, the blue line (left axis) plots the topic-coherence score, which measures the semantic similarity of the words within each topic, and the red line (right axis) plots the in-sample per-word perplexity, which measures model fit. Most of the improvement in both criteria occurs by $K = 50$ (dotted line): coherence rises from 0.38 at $K = 10$ to 0.46 and is flat thereafter, while perplexity declines steeply up to $K = 50$ and only marginally beyond. Hyperparameters are set to $\alpha = 0.05$ and $\eta = 0.02$.

Appendix Tables

Table B.1: Variance decomposition of TGR expectations (OLS) This table presents the explained variation in OLS regressions of TGR expectations on brokerage house, analyst, and firm fixed effects. Panel A replicates the specification in Column (1) of Panel A from Table 4 (i.e., full sample analysis, excluding firms, analysts, and brokerage houses with fewer than 2 observations in total, as their respective fixed effects fully account for the variation in TGR expectations). Panel B excludes brokerage houses with fewer than 2 observations, firms with fewer than 10 observations, and analysts with fewer than 50 observations, replicating the specification in Column (4) of Panel B in Table 4. Both panels also exclude brokerage houses that are not associated with equity analysts who switched employers during the sample period (2000–2023), as it is not possible to separate the effects of these brokerage houses from those of their analysts (Abowd, Kramarz, and Margolis, 1999). Variable definitions appear in Appendix B.

Panel A: Column 1 of Panel A in Table 4	Terminal growth rate$_{i,j,t}$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Brokerage house FE	Yes	No	No	Yes	Yes	No	Yes
Analyst FE	No	Yes	No	Yes	No	Yes	Yes
Firm FE	No	No	Yes	No	Yes	Yes	Yes
Observations	36,200	36,009	36,026	36,009	36,026	35,749	35,749
R^2	0.02	0.47	0.45	0.48	0.47	0.65	0.65
Adjusted R^2	0.02	0.42	0.36	0.42	0.38	0.55	0.55
Panel B: Column 4 of Panel B in Table 4	Terminal growth rate$_{i,j,t}$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Brokerage house FE	Yes	No	No	Yes	Yes	No	Yes
Analyst FE	No	Yes	No	Yes	No	Yes	Yes
Firm FE	No	No	Yes	No	Yes	Yes	Yes
Observations	5,541	5,541	5,414	5,541	5,414	5,414	5,414
R^2	0.12	0.42	0.55	0.43	0.59	0.65	0.66
Adjusted R^2	0.12	0.41	0.49	0.41	0.52	0.59	0.60

Table B.2: Accounting for TGR horizons This table presents the explained variation in OLS regressions of TGR expectations on brokerage house, analyst, firm, and TGR horizon fixed effects, for the subsample for which the information on TGR horizon is available. The TGR horizon denotes the forecast year at which terminal growth begins in the analyst's reported valuation model. Variable definitions appear in Appendix B.

Dependent Variable	Terminal growth rate $_{i,j,t}$				
	(1)	(2)	(3)	(4)	(5)
Brokerage house FE	Yes	Yes	Yes	Yes	Yes
Firm FE	No	Yes	Yes	Yes	Yes
Horizon FE	No	No	Yes	No	Yes
Analyst FE	No	No	No	Yes	Yes
Observations	4,910	4,910	4,910	4,910	4,910
R^2	0.07	0.62	0.63	0.77	0.77

Table B.3: Dispersion in TGR expectations This table presents the results of the dispersion of TGR expectations over firms' life cycle. The unit of observation is at the firm i, analyst j, and forecast year t level. The sample period is 2000–2023. The dependent variable is the squared difference between an analyst's TGR belief and the industry (k) consensus for the firm's age group in forecast year t. The independent variables are indicator variables equal to one if the firm's age falls within the specified range and zero otherwise. Variable definitions appear in Appendix B. The regressions are estimated using ordinary least squares, and the standard errors (in parentheses) are heteroskedastic-consistent and clustered at the firm level. Significance levels are shown as follows: * = 10%, ** = 5%, *** = 1%.

Dispersion in TGR Expectations	$(TGR_{i,j,t} - \overline{TGR}_{k,t}^{\text{Age group}})^2 / \lvert\overline{TGR}_{k,t}\rvert$		
	(1)	(2)	(3)
5 <= Firm age < $10_{i,j,t}$	-0.01	-0.02	-0.02
	(0.02)	(0.02)	(0.03)
10 <= Firm age < $15_{i,j,t}$	-0.01	-0.02	-0.03
	(0.03)	(0.03)	(0.03)
15 <= Firm age < $20_{i,j,t}$	-0.06*	-0.06**	-0.08**
	(0.03)	(0.03)	(0.03)
20 <= Firm age < $25_{i,j,t}$	-0.05	-0.06*	-0.09**
	(0.03)	(0.03)	(0.03)
25 <= Firm age < $30_{i,j,t}$	-0.04	-0.06	-0.08**
	(0.04)	(0.04)	(0.04)
30 <= Firm age < $35_{i,j,t}$	-0.18***	-0.22***	-0.26***
	(0.04)	(0.04)	(0.04)
35 <= Firm age < $40_{i,j,t}$	-0.25***	-0.26***	-0.34***
	(0.04)	(0.04)	(0.04)
40 <= Firm age < $45_{i,j,t}$	-0.29***	-0.28***	-0.34***
	(0.05)	(0.05)	(0.05)
45 <= Firm age < $50_{i,j,t}$	-0.41***	-0.39***	-0.46***
	(0.04)	(0.05)	(0.05)
50 <= Firm age < $55_{i,j,t}$	-0.36***	-0.35***	-0.38***
	(0.05)	(0.06)	(0.06)
55 <= Firm age < $60_{i,j,t}$	-0.45***	-0.47***	-0.55***
	(0.04)	(0.05)	(0.06)
60 <= Firm age < $65_{i,j,t}$	-0.24***	-0.30***	-0.37***
	(0.09)	(0.09)	(0.09)
65 <= Firm age < $70_{i,j,t}$	-0.50***	-0.49***	-0.54***
	(0.06)	(0.04)	(0.05)
70 <= Firm age < $75_{i,j,t}$	-0.26***	-0.33***	-0.43***
	(0.10)	(0.09)	(0.10)
75 <= Firm age < $80_{i,j,t}$	-0.25***	-0.32***	-0.21
	(0.09)	(0.11)	(0.15)
80 <= Firm age < $85_{i,j,t}$	-0.42***	-0.46***	-0.51***
	(0.09)	(0.13)	(0.15)
85 <= Firm age < $90_{i,j,t}$	-0.50***	-0.41***	-0.43***
	(0.07)	(0.08)	(0.09)
90 <= Firm age < $95_{i,j,t}$	-0.29**	-0.36***	-0.41***
	(0.14)	(0.12)	(0.12)
95 <= Firm age < $100_{i,j,t}$	-0.58***	-0.56***	-0.58***
	(0.03)	(0.05)	(0.05)
Firm $\text{Size}_{i,t}$			0.03***
			(0.01)
$\text{TGR}_{i,j,t}$			0.03**
			(0.01)
Constant	0.66***	0.67***	0.01
	(0.02)	(0.02)	(0.12)
Year*Industry FE	No	Yes	Yes
H_0: all age-group coefficients equal, joint test (p-value)	0.00	0.00	0.00
Observations	23,844	23,822	20,941
F Statistics	37.31	20.16	17.09
R^2	0.01	0.08	0.09

Table B.4: **Variance decomposition of TGR by objective predictability of distant growth.** This table repeats the life-cycle decomposition of Table 6 with firm-years sorted into terciles of objective unpredictability, $U_{i,t}$, in place of firm age. $U_{i,t}$ is the expected squared out-of-sample error of a prediction of realized earnings growth from $t+3$ to $t+8$ from public information available at t (the IBES consensus LTG, ln firm age, size, beta, ten-year GDP growth in the firm's country, historical earnings growth from $t-8$ to $t-3$, default risk, and industry and country dummies), estimated leave-one-year-out so that every prediction is out of sample; a second regression of the log squared out-of-sample error on the same information gives $U_{i,t}$ for every firm-year. Within each tercile the sample follows Table 6: firms, analysts, and brokerages with at least two reports in the tercile. Panel A reports the number of reports, firms, and analysts, the variance of TGR, and the average, across firm-years covered by at least two analysts, of the standard deviation of TGR across the analysts covering the same firm in the same year. Panel B reports the sequential ANOVA of Table 6, entering country, industry, brokerage, and analyst fixed effects in that order: the R^2, the shares of explained variation attributable to country and industry factors and to analyst effects, and the analyst component in squared percentage points (sequential sum of squares divided by $N-1$). Because the prediction model includes industry and country dummies, terciles of U differ in country composition; Panel C therefore repeats Panel B with terciles formed by ranking U within country × industry cells (within country for cells with fewer than six firm-years), so that every tercile has the same country and industry mix by construction.

	Tercile of objective unpredictability of distant growth, U		
	(1) Most predictable	(2)	(3) Least predictable
Panel A: Sample and dispersion			
Reports	6,911	6,504	5,271
Firms	963	1,246	1,370
Analysts	1,176	1,275	1,095
Variance of TGR (pp^2)	1.505	1.581	1.715
Within-firm-year SD of TGR (pp)	0.695	0.783	0.798
Variance decomposition			
R^2	0.590	0.569	0.618
Country + industry share of explained variation (%)	23.7	19.4	13.6
Analyst share of explained variation (%)	69.3	75.6	81.1
Country + industry share of total variation (%)	14.0	11.1	8.4
Analyst share of total variation (%)	40.9	43.0	50.1
Analyst component (pp^2)	0.615	0.679	0.859
Panel C: Terciles formed within country × industry			
Reports	6,500	6,416	5,298
Variance of TGR (pp^2)	1.567	1.577	1.680
Within-firm-year SD of TGR (pp)	0.727	0.780	0.748
R^2	0.576	0.582	0.635
Country + industry share of explained variation (%)	23.1	17.3	16.0
Analyst share of explained variation (%)	70.0	77.2	79.5
Country + industry share of total variation (%)	13.3	10.1	10.2
Analyst share of total variation (%)	40.3	45.0	50.5
Analyst component (pp^2)	0.630	0.708	0.849

Table B.5: Decomposition of analyst fixed effects (OLS) This table complements the sequential ANOVA decomposition of Panel B in Table 7 with an order-free assessment of how much of the variation in analyst fixed effects each set of observable characteristics explains. The dependent variable is the analyst fixed effect estimated in Panel A of Table 7 from a regression of TGR expectations on firm, brokerage-house, and analyst fixed effects; the unit of observation is the analyst. The table reports the R^2 of OLS regressions of this effect on each set of characteristics alone (columns (1) to (5)), on all sets except analyst country (column (6)), and on all sets (column (7)), so that the difference between columns (6) and (7) gives the incremental explanatory power of analyst country. Gender is a binary variable equal to 1 if the analyst is female and 0 otherwise. Education is a categorical variable indicating a master's degree, an MBA, or neither. Race is a categorical variable indicating White, Asian, Hispanic, and other races. Birth cohort is a set of birth year indicators. Analyst country is a set of indicators for the country from which the analyst reports most often. Panel A uses the sample of column (1) of Table 7: firms, analysts, and brokerage houses with fewer than two observations are excluded, as their respective fixed effects fully explain the associated variation in TGR expectations. Panel B uses the sample of column (6) of Table 7, which raises this requirement to five observations and retains only the foreign reports on firms with some foreign coverage. Both panels also exclude brokerage houses not associated with any analyst who switched employers during the sample period (2000–2023), as the effects of these brokerage houses cannot be separated from those of their analysts (Abowd, Kramarz, and Margolis, 1999), and require that analysts from at least two countries cover each firm country. Variable definitions appear in Appendix Appendix B.

Panel A: Column 1 of Panel B in 7	Analyst FE_j						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Gender FE	Yes	No	No	No	No	Yes	Yes
Education FE	No	Yes	No	No	No	Yes	Yes
Race FE	No	No	Yes	No	No	Yes	Yes
Birth cohort FE	No	No	No	Yes	No	Yes	Yes
Analyst Country FE	No	No	No	No	Yes	No	Yes
Observations	955	955	955	955	955	955	955
R^2	0.00	0.00	0.00	0.08	0.10	0.08	0.17
Panel B: Column 6 of Panel B in 7	Analyst FE_j						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Gender FE	Yes	No	No	No	No	Yes	Yes
Education FE	No	Yes	No	No	No	Yes	Yes
Race FE	No	No	Yes	No	No	Yes	Yes
Birth cohort FE	No	No	No	Yes	No	Yes	Yes
Analyst Country FE	No	No	No	No	Yes	No	Yes
Observations	291	291	291	291	291	291	291
R^2	0.00	0.01	0.01	0.16	0.11	0.18	0.30

Table B.6: Representative words from LDA topics This table reports the key words representing LDA topics. Words are selected based on reversed frequency weights, where a word's weight in a given topic is scaled by its relative frequency in the texts, similar to the strategy in Bybee, Kelly, and Su (2023).

Category	Topic Name	Word List
Geography	Asia	rmb, thailand, hong_kong, korea, china, overseas, singapore, india, chinese, domestic
Geography	Europe	spanish, italian, spain, italy, france, french, russian, germany, german, russia
Geography	Oceania	australian, new_zealand, australia, half, earning, forecast, period, earn, result
Geography	South-America	mexican, brazilian, argentina, brazil, chile, mexico, latam, inflation, country
Industry	Health Care	hospital, implant, medicare, pharmacy, diagnostic, medical, care, healthcare, reimbursement, procedure
Industry	Oil and Gas (Transformation)	refining, petrochemical, refinery, feedstock, crude_oil, gasoline, crude, bbl, downstream, oil
Industry	Oil and Gas (Production)	rig, boe, drilling, exploration, field, discovery, reserve, natural_gas, oil_gas, oil
Industry	Manufacturing	equipment, industrial, manufacturing, manufacturer, chemical, system, supplier, component, technology, application
Industry	Entertainment	programming, broadcast, newspaper, television, film, music, radio, audience, entertainment, advertising
Industry	Resources	ethanol, pulp, farmer, sugar, crop, grain, cigarette, corn, seed, fertilizer
Industry	Automobile	battery, dealer, car, truck, vehicle, auto, engine, automotive, man, unit
Industry	Banking	loan, lending, bank, mortgage, banking, insurance, deposit, fund, bond, roe
Industry	Carrier	satellite, tower, fiber, cable, video, carrier, network, voice, speed, broadband
Industry	Communication	arpu, wireline, mobile, subscriber, telecom, churn, broadband, wireless, sub, operator
Industry	Construction	non_residential, toll_road, cement, housing, railway, concession, construction, road, contractor, building
Industry	Consumer goods (Clothing)	apparel, fashion, luxury, brand, wholesale, sport, distributor, channel, inventory, retail
Industry	Consumer goods (Food)	beverage, spirit, tobacco, beer, wine, category, innovation, brand, food, consumer
Industry	Energy	renewable, grid, mwh, hydro, solar, electricity, power_plant, wind, nuclear, utility
Industry	Gaming	gaming, search, game, online, commerce, website, internet, user, platform, marketing
Industry	IT	cloud, software, sap, vendor, solution, enterprise, subscription, application, license, platform
Industry	Mining (Metal Other)	gold, mineral, zinc, south_african, copper, ore, nickel, mining, south_africa, grade
Industry	Mining (Steel)	potash, aluminium, steel, iron_ore, aluminum, coal, tonne, ton, import, producer
Industry	Pharmaceutical (R&D)	generic, pharma, pharmaceutical, patent, drug, fda, royalty, launch, franchise, pipeline
Industry	Pharmaceutical (Clinical trial)	phase_trial, dose, efficacy, therapy, cancer, disease, patient, trial, clinical, study
Industry	Real Estate	reit, tenant, occupancy, property, hotel, rent, rental, real_estate, land, nav
Industry	Retail (Other)	supermarket, store, restaurant, department_store, retailer, format, food, retail, chain, comp
Industry	Semi-Conductor	dram, chip, semiconductor, panel, handset, shipment, tech, device, design, technology
Industry	Shipping	container, vessel, parcel, port, mail, ship, rail, tanker, freight, shipping
Industry	Transportation	airline, aircraft, cruise, passenger, airport, aviation, flight, defense, travel, toll
Macroeconomics	Legal Risk	court, litigation, settlement, legal, card, claim, merchant, affiliate, payment, information
Macroeconomics	Economics Cycles	economy, gdp, index, economic, cycle, recovery, global, weak, sector, decline
Macroeconomics	Regulation	reform, proposal, regulator, regulation, regulatory, government, tax, review, state, policy

Continued on next page.

Continued from previous page.

Category	Topic Name	Word List
Operations	Outsourcing	outsourcing, client, service, contract, employee, provider, revenue, business, segment, acquisition
Operations	Demand	volume, pricing, raw_material, demand, capacity, supply, pressure, increase, recovery, cost
Operations	Expenditure and Disbursement	dps, dividend, debt, buyback, capex, dividend_yield, cash, fcf, balance_sheet, ebitda
Operations	Management	officer, education, student, school, director, executive, president, esg, cfo, ceo
Operations	Merge and Acquisition	merger, bid, synergy, deal, stake, transaction, shareholder, offer, acquisition
Operations	Organization	disposal, restructuring, division, group, trading, european, profit
Operations	Project Management	waste, project, water, backlog, delay, pipeline, facility, contract, financing, funding
Valuation	Valuation facts	fact, different, reason, important, big, time, long, case, point, issue
Valuation	Accounting Norms and Model Selection	ambiguity_local, modelware_separation, accounting_regulation, value_distortion, financing_complete, fifo_basis, inventory_lifo, eps_adjust, modelware_proprietary, interest_objectivity
Valuation	Free Cash Flow	approximately, roughly, free_cash, flow, facility, cash_flow, company, capital, distribution, acquisition
Valuation	Operating Perfomance	sek, organic, ebit, ebita, currency, sale, europe, margin, effect, growth
Valuation	Profitability	profit, expense, income, loss, operating, gross, charge, net, gain, tax
Valuation	Firm growth	covid, headwind, bps, incremental, leverage, fcf, expansion, organic, upside, margin
Valuation	Recommendation	outperform, neutral, catalyst, recommendation, stock, momentum, buy, rating, earning, earn
Valuation	Guidance	eur, usd, fcf, ebit, slightly, guidance, positive, estimate, good, development
Valuation	Consensus	guidance, street, consensus, estimate, eps, prior, revenue, slightly, expectation, ahead
Valuation	Relative Valuation	multiple, fair_value, historical, peer, analysis, average, discount, relative, scenario, value
Valuation	Risk	beta, wacc, risk, terminal, downside, free, dcf, investment_thesis, hold, unchanged

Table B.8: TGR disagreement and textual discussions (Alternative topic model with 80 topics) This table presents the relationship between differences in textual discussion and disagreement in TGR expectations between pairs of analysts, using an alternative LDA model with 80 topics. The unit of observation is at the analyst a, analyst b, firm i, and forecast year t level. The sample period is 2000–2023. The dependent variable, $Disagreement_{a,b,i,t}$, is defined as the squared difference in TGR estimates between analyst a's and analyst b's TGR beliefs: $(TGR_{a,i,t} - TGR_{b,i,t})^2$. The key independent variables of interest are the Jaccard and Cosine similarity scores, which capture the similarity of the analysts' textual discussion on the extensive and intensive margins, respectively. Variable definitions appear in Appendix B. The regressions are estimated using ordinary least squares, and the standard errors (in parentheses) are heteroskedastic-consistent and clustered at the analyst-pair level. Significance levels are shown as follows: * = 10%, ** = 5%, *** = 1%.

Dependent variable	$Disagreement_{a,b,i,t}$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Jaccard $^{80}_{a,b,i,t}$	-0.19	-0.33	-0.30				0.02
	(0.28)	(0.28)	(0.28)				(0.24)
Cosine $^{80}_{a,b,i,t}$				-0.53***	-0.49**	-0.37*	-0.27
				(0.20)	(0.20)	(0.20)	(0.17)
Year FE	No	Yes	No	No	Yes	No	No
Industry FE	No	Yes	No	No	Yes	No	No
Industry*Year FE	No	No	Yes	No	No	Yes	Yes
Analyst pair FE	No	No	No	No	No	No	Yes
Firm FE	No	No	No	No	No	No	Yes
Brokerage pair FE	No	No	No	No	No	No	Yes
Observations	23,466	23,433	23,400	23,463	23,430	23,397	19,750
F Statistics	0.46	1.43	1.19	6.70	6.01	3.56	1.25
R^2	0.00	0.02	0.07	0.00	0.02	0.07	0.74

Table B.9: Split-sample persistence of analyst effects This table tests whether the analyst fixed effects reflect persistent differences across analysts rather than sampling error. For each random split, every analyst's coverage is divided into two non-overlapping halves, and the analyst effect is estimated separately in each half from a regression of TGR on analyst, brokerage house, firm, and firm-country×industry×year fixed effects. Because estimation error is independent across halves, the covariance between the two estimates isolates the persistent component of the analyst effect. In the relationship split (columns (1)–(2)), entire analyst–firm relationships are assigned to one half, so the two halves share no firm within an analyst and the covariance can only arise from the analyst carrying the same level across different firms; each analyst must cover at least two distinct firms, with at least five reports, in each half. In the analyst-firm-year split (columns (3)–(4)), analyst-firm-year observations are assigned to the halves at random, with at least five reports per half. Analyst effects are demeaned within each half and weighted by the analyst's number of reports (the unweighted correlation is reported for comparison). The cross-half predictive slope is the covariance divided by the variance of the first-half estimate, i.e., the coefficient from regressing the second-half effect on the first-half effect. The reliability of the full-sample analyst effect is the share of the variance of the effect estimated on the full sample that is persistent, cov/[cov + (noise variance)/4]; it is the shrinkage factor applied to the analyst effects in Table 10. Each design is repeated over 100 random splits; the table reports the median and the 25th and 75th percentiles across splits, which describe sensitivity to the random partition and are not confidence intervals. Fixed effects are identified only up to a constant within each connected component, and these normalizations are independent across halves, which attenuates the cross-half covariance; the reported persistence is therefore conservative.

	Relationship split		Analyst-firm-year split	
	Median	[25th, 75th]	Median	[25th, 75th]
	(1)	(2)	(3)	(4)
Persistent component of analyst effects				
Covariance of effects across halves (pp^2)	0.402	[0.373, 0.436]	0.381	[0.358, 0.397]
SD of the persistent component (pp)	0.634	[0.611, 0.660]	0.618	[0.598, 0.630]
Agreement between the two halves				
Correlation across halves, weighted	0.418	[0.386, 0.460]	0.532	[0.491, 0.572]
Correlation across halves, unweighted	0.387	[0.355, 0.432]	0.506	[0.471, 0.539]
Cross-half predictive slope	0.421	[0.367, 0.468]	0.532	[0.472, 0.590]
Implication for the full-sample effects				
Reliability of the full-sample analyst effect	0.589	[0.554, 0.629]	0.692	[0.658, 0.726]
Sample				
Analysts with effects in both halves	635	[626, 640]	700	[693, 705]
Completed random splits	100		100	
Splits with negative covariance (%)	0.0		0.0	

Table B.10: Robustness of the TGR variance decomposition This table re-estimates the sequential variance decomposition of TGR expectations in column (1) of Table 4 on three alternative samples and specifications. The unit of observation is at the firm i, analyst j, and forecast year t level; the sample period is 2000–2023. In each column, TGR expectations are regressed on brokerage house, analyst, and firm fixed effects, and the table reports, for each set of fixed effects, the share of the explained variation (R^2) attributed to it, the share of total variation it explains, and its degrees of freedom, with the fixed effects entered in the order listed. Column (1) reproduces the baseline: the sample excludes firms, analysts, and brokerage houses with fewer than two observations and brokerage houses not associated with any analyst who switched employers during the sample period, as the effects of these brokerage houses cannot be separated from those of their analysts (Abowd, Kramarz, and Margolis, 1999). Column (2) restricts the sample further to the analysts who themselves switched brokerage house, so that every analyst effect is identified across employers rather than only within a connected set of houses. Column (3) uses the baseline sample but replaces the brokerage-house fixed effects with brokerage house × industry × year fixed effects, which absorb any house-wide view of an industry in a given year, such as a common macroeconomic outlook or a house valuation template, before the analyst and firm effects are added. Column (4) drops the reports whose TGR is a multiple of 0.5 percentage point, to verify that the share attributed to analysts does not reflect differences in analysts' conventions of reporting round numbers. Variable definitions appear in Appendix Appendix B.

	Terminal growth rate$_{i,j,t}$			
	Baseline (Table 4, col. 1) (1)	Switching analysts only (2)	Brokerage × industry × year FE (3)	Non-round TGRs only (4)
Brokerage house FE (brokerage × industry × year in (3))				
Share of explained R^2 (%)	3	4	34	15
Share of total variation (%)	2	3	25	14
Degrees of freedom	39	36	4,077	38
Analyst FE				
Share of explained R^2 (%)	69	56	44	56
Share of total variation (%)	46	36	33	51
Degrees of freedom	3,493	378	3,310	1,072
Firm FE				
Share of explained R^2 (%)	27	40	22	28
Share of total variation (%)	18	25	16	26
Degrees of freedom	5,029	1,407	4,920	1,302
Total R^2	0.66	0.64	0.74	0.90
Observations	36,200	6,651	34,581	3,828