Analyst Belief Overreaction (LTG) and Stock Returns: Technology versus Non-Technology in U.S. Equities

by

Mohammad Sheikhi

An Extended Essay Submitted in Partial Fulfillment of the Requirements for the Degree of

MASTER OF ARTS

in the Department of Economics

Dr. Ke Xu, Supervisor (Department of Economics)

We acknowledge and respect the Lkŋn (Songhees and Xsepsm/Esquimalt) Peoples on whose territory the university stands, and the Lkŋn and WSÁNEĆ Peoples whose historical relationships with the land continue to this day.

Abstract

This essay reviews evidence on whether analysts' long-term earnings growth (LTG) beliefs predict subsequent stock returns via belief overreaction and whether this predictability differs between technology and non-technology firms in the United States (December 1982–July 2025). Synthesizing work on LTG levels, revisions, dispersion, and related proxies, the evidence indicates that elevated LTG or optimism corresponds to lower future returns, consistent with overreaction and subsequent correction. Building on Bordalo et al. (2022), whose firm-level tests provide strong evidence of belief-driven return predictability, and motivated by sectoral differences, the analysis also tests for heterogeneity between technology and non-technology firms. Although theory suggests potential amplification in technology, the IV estimates do not indicate a statistically reliable difference in the relationship for technology versus non-technology.

Contents

List of Tables

List of Figures

1 Introduction

Understanding how expectations about distant cash flows are formed, transmitted into market prices, and ultimately corrected is central to asset pricing. Analysts' long-term earnings growth (LTG) forecasts—as compiled by I/B/E/S [1]—offer one of the only systematic, high-frequency windows into such beliefs. By construction, LTG is a multi-year (roughly three- to five-year) expectation of average annual EPS growth and, therefore, a natural proxy for beliefs about the persistent component of fundamentals. Because a large share of equity value for growth firms reflects distant cash flows, errors in LTG should have first-order pricing consequences. This channel matches behavioral-finance models in which investors extrapolate salient recent growth into the distant future, elevating LTG and prices above fundamentals and generating subsequent mean reversion as information is revealed (Bordalo et al. 2022; La Porta 1996).

Historical episodes underline this mechanism. During the late-1990s technology boom, analysts issued very optimistic long-horizon expectations for internet and telecom firms, often accompanied by aggressive price targets. As the cycle turned, realized outcomes disappointed and prices corrected sharply. See Lamont and Thaler (2003) for evidence from technology carve-outs. Although such episodes are extreme, they illustrate how sectoral characteristics influence the relationship between belief errors and returns. This essay addresses two questions: (i) whether long-horizon earnings forecast errors predict discounted five-year returns at the firm level and (ii) whether this relationship differs systematically between technology and non-technology firms. From a research-design perspective, splitting the cross-section offers two tests. Technology is defined using GICS industry-group codes 4510, 4520, and 4530; all other firms are non-technology (see Section 3).

The foregoing literature primarily studies market-wide relations (Lakonishok et al. 1994; Baker and Wurgler 2006; Stambaugh et al. 2012). This essay contributes a sectoral perspective: do LTG-related return-predictability patterns vary systemically between technology and

[1] I/B/E/S = Institutional Brokers' Estimate System (provided via the LSEG data platform).

non-technology firms? A priori, technology is a natural locus for larger belief cycles. First, fundamentals are more intangible and option-like. A larger share of value is tied to uncertain projects whose payoffs lie far in the future, making long-run beliefs both harder to form and more consequential for valuation. Second, technology enjoys unusually high investor attention. Sentiment tends to concentrate in hard-to-value securities with option-like payoffs and high dispersion, and mispricing pressures are strongest when attention and disagreement are elevated (Baker and Wurgler 2006). Third, limits to arbitrage are often binding in speculative runs, allowing deviations to persist (Shleifer and Vishny 1997). Each channel amplifies the impact of extrapolative beliefs: when a compelling narrative meets uncertainty and frictions, overreaction is more severe and more persistent.

The belief measures studied in the literature inform the empirical choices. The analysis considers (i) LTG levels, (ii) LTG revisions, and (iii) disagreement as core constructs. Levels and revisions capture extrapolation in the mean belief; disagreement captures the breadth of views and, with short-sale frictions, the tendency for optimistic beliefs to dominate prices. Portfolio sorts on these constructs yield transparent economic magnitudes, while panel regressions with fixed effects provide statistical control. For identification at a long horizon, following Bordalo et al. (2022), long-run forecast errors are instrumented with lags and changes in LTG. Inference accounts for dependence across time and firms; The model uses Driscoll–Kraay (DK) standard errors with a long bandwidth appropriate for long-horizon panels.

In a firm–time fixed-effects IV using LTG's 12-month *lag* and *change* as instruments, the first stage is strong and the IV slope of discounted five-year returns on five-year forecast errors is positive and economically meaningful. This is consistent with belief-correction dynamics: when long-run beliefs are overly optimistic, expected returns are temporarily low and realized returns are subsequently higher as errors dissipate. In pooled interaction IVs allowing technology-specific slopes, point estimates are steeper for technology but are not statistically significant; standard Wald tests do not reject equal coefficients for technology vs non-technology companies.

In sum, decades of research show that LTG and related belief measures contain information about the cross-section of expected returns that is difficult to reconcile with purely rational risk premia. The diagnostic-expectations framework rationalizes these facts, and firm-level IV evidence ties long-run belief errors to subsequent returns. The open question is whether, and by how much, these patterns are stronger in technology. Ex ante, the expectation is that the relationship between beliefs and prices should be steeper in technology, given higher uncertainty, greater attention, and tighter arbitrage frictions. Establishing this heterogeneity sharpens the understanding of when belief overreaction moves markets the most and guides both theory and practice toward the environments in which mispricing is most likely to be found.

This essay proceeds as follows: Section 2 reviews related work and positions the contribution. Section 3 describes data, variable construction, and sample formation. Section 4 details the econometric design. Section 5 reports the baseline IV and the technology–non-technology interaction results. The Tables and Figures sections provide results, robustness checks, and figures. The appendix represents operational variable definitions.

2 Related Literature

A seminal insight is that analyst expectations are not merely noisy but systematically extrapolative. La Porta (1996) sorts firms on consensus LTG and documents that portfolios with the most optimistic long-run growth forecasts subsequently underperform those with the most pessimistic forecasts by economically large margins. The spread persists after standard risk adjustments and coincides with later earnings disappointment by the initially favored firms. The implication is straightforward: when beliefs are overly optimistic, prices embed overvaluation that reverts as reality catches up. Subsequent work, most notably Chan et al. (2003), shows that extraordinary growth is rarely sustained and that analysts, on average, overestimate persistence. This combination—rare persistent growth and extrapolative forecasting—naturally generates predictable reversals in prices.

A theoretical structure for such findings is provided by the diagnostic-expectations paradigm of Bordalo et al. (2019). Agents overweight recent, salient signals when projecting the future; applied to corporate fundamentals, recent earnings strength is overweighted in LTG, generating too-optimistic long-run forecasts. Prices that rationally reflect those beliefs will be too high, and expected returns too low, until the error dissipates. Bordalo et al. (2022) push this logic to firm-level identification: they construct long-horizon forecast errors (realized long-run EPS growth minus contemporaneous LTG) and instrument those errors using lagged LTG levels and changes, alongside firm and year fixed effects. The IV estimates indicate that optimism in long-run beliefs predicts lower five-year returns, and first-stage robustness checks are strong. Their conclusion is that belief overreaction is a quantitatively important driver of asset-pricing "puzzles," unifying a range of anomalies within a single expectations framework.

Importantly, belief overreaction appears across multiple analyst-based proxies. Forecast dispersion—cross-analyst disagreement—is negatively related to subsequent returns, as in Diether et al. (2002). This is natural under short-sale constraints and heterogeneous priors: optimistic views are overrepresented in the marginal price, inflating valuations. Analyst *recommendations* and *target prices* embed LTG assumptions, and Bradshaw (2004) shows that analysts lean heavily on growth inputs when issuing ratings; the most optimistic ratings subsequently underperform on average. These auxiliary facts reinforce a single message: when the sell-side narrative is especially upbeat, expected returns tend to be low.

A competing risk-based interpretation posits that LTG proxies for exposures associated with lower discount rates. Growth firms often have high valuations and low distress risk; perhaps low future returns reflect low risk premia rather than mispricing. In principle, long-duration cash-flow profiles could also make growth firms sensitive to discount-rate shocks. Yet three considerations weigh against a pure risk story. First, the direction of the forecast-dispersion premium (high disagreement $\Rightarrow$ low returns) runs counter to standard risk logic (Diether et al. 2002). Second, the predictive power of LTG for returns is accompanied by earnings disappointments for the high-LTG firms and favorable surprises for low-LTG firms (La Porta 1996), which matches an expectations-correction channel, not the payoffs to bearing systematic risk. Third, IV-based evidence at the firm level links *forecast errors*—not simply long-run levels—to subsequent returns (Bordalo et al. 2022). If the signal were purely risk, instrumented expectation errors would not systematically forecast returns once fixed effects and controls are included. While multifactor models can absorb some average spreads by construction, the weight of evidence points to belief dynamics as first-order.

Two additional literatures complement the analyst-belief evidence. First, classic contrarian results emphasize that investors over-extrapolate past performance; value strategies earn premia in part because expectations for "glamour" firms are too optimistic (Lakonishok et al. 1994). Analyst LTG is a direct, higher-frequency belief proxy for the same mechanism. Second, the sentiment literature links broad investor mood to anomalies: mispricing is stronger when sentiment is elevated, particularly in speculative, hard-to-arbitrage stocks (Baker and Wurgler 2006; Stambaugh et al. 2012). Technology is frequently overrepresented in those categories, providing a macro-state in which LTG effects should be more pronounced. Together, these perspectives point to an interaction: belief-driven return predictability should vary with both *sector* (technology versus non-technology) and *sentiment state* (high versus low).

This essay contributes in two ways: it replicates the firm-level belief-overreaction result in an updated U.S. sample and assesses sectoral heterogeneity by comparing technology with non-technology firms. The main effect is replicated in an updated U.S. firm sample (December 1982–July 2025): when analysts' long-run growth beliefs overshoot what firms ultimately deliver, five-year returns tend to be higher afterward, consistent with a gradual correction of overreaction. Also, the analysis examines whether this pattern is stronger for technology companies. Using a clear technology–non-technology split based on GICS groups, the relationship is larger for technology, but the difference is not statistically significant.

Taken together, the paper reaffirms the belief-overreaction channel and suggests that any technology amplification, if present, is modest.

3 Variable Construction and Data

This section introduces the data, describes how the estimation sample is formed, and defines the variables used throughout the paper. The presentation is intentionally narrative and selective: operational minutiae (e.g., field-by-field mappings and implementation details) are recorded in the Appendix so that Section 3 remains readable and focused on economic content. Brief motivations accompany each definition; more extensive conceptual discussion is in Section 2 (Related Literature).

The analysis uses a monthly panel of U.S. common stocks with valid Refinitiv instrument identifiers. For each firm i and calendar month t between December 1982 and July 2025[2], analysts' long-term earnings growth (LTG) expectations, trailing-twelve-month earnings per share (EPS), standard fundamentals (market capitalization, book value per share, price), and industry classifications are observed. The monthly frequency aligns naturally with long-horizon compounding and with the timing discipline used to construct outcomes and instruments below.

Industry membership is measured using the *GICS Industry Group* taxonomy. Firms are classified as *technology* if their group is Software & Services, Technology Hardware & Equipment, or Semiconductors & Semiconductor Equipment (codes 4510, 4520, 4530); all other firms are labeled *non-technology.* This time-invariant indicator underlies the heterogeneity analysis and the matched-sample design described next.

The analysis begins from the full monthly panel and applies a small number of standard filters to obtain the estimation sample used in Sections 4–5. First, the sample retains firm–months with non-missing LTG and total return. Second, because long-run EPS growth

[2]Because $r^{(5)}$ and $e^{(5)}$ use future outcomes, the last month with non-missing values is effectively about five years before the raw end date.

is computed in logs, strictly positive EPS is required at the relevant endpoints. Third, to mitigate the influence of extreme outliers and data glitches, key variables are winsorized at the 1st/99th percentiles separately for each calendar month.[3] Fourth, where indicated in the results, variables are standardized by their full-sample standard deviation for interpretability; when this is done, the instruments are formed from the standardized LTG series for internal consistency. Finally, for the technology heterogeneity tests the panel is restricted to the propensity-score matched universe (1:1 nearest neighbor on time-averaged size, valuation, return, and leverage), ensuring that technology and non-technology firms are matched on key observables.[4]

As a compact check of the firm propensity matching, Table 2 compares technology and non-technology firms' main characteristics before and after the propensity-score match: time-averaged size (market capitalization), valuation (book-to-price), average monthly return, and leverage (debt/assets). For each variable, the number of available observations and the cross-sectional mean by sector are reported, together with the percent difference (non-technology relative to technology). Counts differ across rows because of missing data. After matching, the large leverage gap shrinks markedly (from about 51% to 4.5%); the valuation gap narrows modestly; average returns remain similar; and the size gap persists, indicating that the propensity-score matching was successful and minimized differences between groups.

This subsection defines the main variables used in the analysis. The analysis requires: (i) monthly total returns and prices; (ii) analysts' LTG (percent, converted to decimals); (iii) TTM EPS for constructing realized long-run growth; (iv) GICS groups for technology status; and (v) baseline covariates used only for matching (size, B/P, prior 12-month average return, leverage).

Let $r_{\log,i,t} \equiv \ln\Big(1 + r_{i,t}\Big)$ denote the monthly log return, where i indexes firms and t calendar months. The five-year outcome is the *discounted* sum of future monthly log returns

[3]Month-wise winsorization is standard in long-horizon panels; it reduces the influence of transient spikes without pooling information across calendar time.

[4]The matching is implemented at the firm level and produces treated–control instrument pairs used to filter the panel; details and robustness checks are reported with the summary statistics.

in annual blocks, with $j \in \{0, \ldots, 4\}$ indexing the five annual blocks, $k \in \{1, \ldots, 12\}$ the months within each block, and $\rho \in (0, 1)$ the constant annual discount factor; $r_{i,t}^{(5)}$ denotes the discounted five-year log return:

$$r_{i,t}^{(5)} = \sum_{j=0}^{4} \rho^j \left(\sum_{k=1}^{12} r_{\log,i,t+12j+k} \right), \qquad \rho \in (0, 1). \tag{1}$$

Discounting recognizes that more distant cash-flow news is less salient for present value. Following Bordalo et al. (2022), annual blocks are weighted using a constant annual discount for this sample and horizon, $\rho \approx 0.978$, matching their long-horizon specification.

Let $\text{EPS}_{i,t}$ denote trailing-twelve-month EPS and $\text{LTG}_{i,t}$ denote the long-term EPS growth expectation (in decimals). To respect information flow, the base EPS is lagged by three months; the numerator looks forward (57 months to align with monthly indexing). The five-year forecast error is

$$e_{i,t}^{(5)} = \underbrace{\frac{\ln(\text{EPS}_{i,t+57}) - \ln(\text{EPS}_{i,t-3})}{5}}_{\text{annualized realized 5y EPS growth}} - \text{LTG}_{i,t}. \tag{2}$$

This construction compares realized long-run earnings to what analysts had expected at t. A three-month lag mitigates look-ahead bias from reporting delays; it is a standard choice in the long-horizon literature.

The forecast error $e_{i,t}^{(5)}$ is measured with noise and may reflect omitted news correlated with returns. A three-year analogue $e_{i,t}^{(3)}$ is constructed identically to $e_{i,t}^{(5)}$ but over a three-year horizon; the corresponding discounted return outcome is $r_{i,t}^{(3)}$. These are used only for robustness (Tables 6 and 7). Following best practice, *LTG-based* instruments dated at t and $t-12$ are used:

$$\text{LTG}_{i,t-12} \text{ (12-month lag of LTG)}, \qquad \Delta\text{LTG}_{i,t} \equiv \text{LTG}_{i,t} - \text{LTG}_{i,t-12}.$$

These terms are predictive of $e_{i,t}^{(5)}$ in the first stage and, after absorbing firm and year fixed effects, are plausibly orthogonal to innovations in $r^{(5)}_{i,t}$. In the pooled interaction IV that allows the slope on $e^{(5)}_{i,t}$ to differ by technology status, the interaction $e^{(5)}_{i,t} \times D^{\text{Tech}}_i$ is instrumented with $\text{LTG}_{i,t-12} \times D^{\text{Tech}}_i$ and $\Delta\text{LTG}_{i,t} \times D^{\text{Tech}}_i$.

Technology status is the time-invariant indicator

$$D^{\text{Tech}}_i = \begin{cases} 1, & \text{if GICS Industry Group} \in \{4510,\ 4520,\ 4530\}, \\ 0, & \text{otherwise.} \end{cases}$$

As a complementary descriptive check focused on the matched universe, Table 1 reports distributional summaries (mean, standard deviation, and selected quantiles) by sector for the variables used in the empirical analysis. This table includes scale and dispersion that underlie the IV estimates and confirms sectoral patterns. Consistent with priors, technology firms exhibit higher LTG levels and volatility, more negative and dispersed long-horizon forecast errors, and slightly higher monthly return volatility; realized long-horizon returns $r^{(3)}$ and $r^{(5)}$ have similar means across sectors but heavier tails in technology. The instruments mirror this: $\text{LTG}_{i,t-12}$ and ΔLTG are, on average, higher and more volatile in technology, supporting first-stage relevance.

To document the macro context for the instruments and outcome, Figure 1 plots cross-sectional means of lagged LTG ($\text{LTG}_{i,t-12}$), the five-year forecast error $e^{(5)}$, and discounted five-year returns $r^{(5)}$. Beliefs and errors are cyclical with pronounced excursions around the late-1990s boom and the Global Financial Crisis, patterns that motivate the LTG-based instruments (lag and change) and the robustness that excludes those windows in Section 5.

4 Methodology

This section presents the econometric design used to quantify the relationship between long-horizon earnings forecast errors and discounted five-year returns and to assess whether this relationship differs between technology and non-technology firms. The outcome $r^{(5)}_{i,t}$ and the forecast error $e_{i,t}^{(5)}$ follow the constructions in Section 3. The structural specification, the instrument set and identification logic, and the estimation and inference procedures are set out below, followed by the technology interaction.

The baseline specification relates discounted five-year log returns to the five-year forecast error within a panel with firm and time controls:

$$r_{i,t}^{(5)} \;=\; \alpha_i \;+\; \lambda_t \;+\; \beta\, e_{i,t}^{(5)} \;+\; \varepsilon_{i,t}, \tag{3}$$

where α_i are firm fixed effects and λ_t are year fixed effects. The regressor $e_{i,t}^{(5)}$ is treated as endogenous to account for measurement error and potential co-movement with omitted news that also affect returns.

Two instruments based on long-term growth expectations (LTG) shift long-run beliefs in ways that are predictive of the forecast error: the 12-month change $\Delta\mathrm{LTG}_{i,t} \equiv \mathrm{LTG}_{i,t} - \mathrm{LTG}_{i,t-12}$ and the 12-month lag $\mathrm{LTG}_{i,t-12}$. Conditional on firm and year fixed effects, these variables are designed to affect discounted five-year returns only through their effect on $e_{i,t}^{(5)}$. Relevance is evaluated in the first stage; exclusion is supported by the timing discipline embedded in $e_{i,t}^{(5)}$—the base EPS is lagged by three months, and the realized component looks forward five years while accounting for time-invariant heterogeneity and common shocks via fixed effects.

The first-stage projection is

$$e_{i,t}^{(5)} \;=\; \alpha_i \;+\; \lambda_t \;+\; \pi_1\, \Delta\mathrm{LTG}_{i,t} \;+\; \pi_2\, \mathrm{LTG}_{i,t-12} \;+\; u_{i,t}. \tag{4}$$

Equation (3) is estimated by 2SLS with $\{\Delta\mathrm{LTG}_{i,t},\, \mathrm{LTG}_{i,t-12}\}$ as excluded instruments.

All variables follow the Section 3 preprocessing pipeline. Variables are winsorized by calendar month at the 1st/99th percentiles. Where variables are standardized for interpretability, the instruments are formed from the standardized LTG series so that the exclusion restriction is preserved under re-scaling. Core specifications include firm and year fixed effects. Inference uses Driscoll–Kraay (DK) standard errors with a long bandwidth (approximately 60 months) appropriate for heteroskedasticity, serial correlation, and cross-sectional dependence in long-horizon panels. The Kleibergen–Paap rk Wald F statistic for first-stage strength is reported. As a robustness check, the model is also estimated at a three-year horizon using the analogous constructs $r^{(3)}$ and $e^{(3)}$; inference uses Driscoll–Kraay standard errors with a 36-month bandwidth (see Tables 6 and 7).

To allow for sectoral differences, the slope on the forecast error is allowed to vary with technology status using the time-invariant indicator D_i^{Tech} (Section 3):

$$r_{i,t}^{(5)} \;=\; \alpha_i + \lambda_t + \beta_1\, e_{i,t}^{(5)} + \beta_2 \left(e_{i,t}^{(5)}\, D_i^{\text{Tech}}\right) + \varepsilon_{i,t}. \tag{5}$$

The coefficient β_1 captures the non-technology slope; the technology slope equals $\beta_1 + \beta_2$. Both $e_{i,t}^{(5)}$ and $e_{i,t}^{(5)} D_i^{\text{Tech}}$ are endogenous. They are instrumented with

$$\mathcal{Z}_{i,t} = \begin{Bmatrix} \Delta\text{LTG}_{i,t},\ \text{LTG}_{i,t-12}, \\ \Delta\text{LTG}_{i,t} D_i^{\text{Tech}},\ \text{LTG}_{i,t-12} D_i^{\text{Tech}} \end{Bmatrix}.$$

The system is estimated by 2SLS with firm and year fixed effects and Driscoll–Kraay (DK) standard errors. Because the system is overidentified, first-stage robustness checks and weak-IV-robust inference are reported. The matched-sample universe introduced in Section 3 is used throughout so that technology and non-technology firms are compared on a matched set of observables under the same fixed-effects structure. The associated reduced-form first stages for the endogenous regressor appear in Eq. (6); it map the instrument and its interactions into $e_{i,t}^{(5)}$ and $e_{i,t}^{(5)} D_i^{\text{Tech}}$.

The first stage for $e_{i,t}^{(5)}$ is:

$$e_{i,t}^{(5)} \;=\; \alpha_i^{(1)} + \lambda_t^{(1)} + \pi_1\, \Delta\text{LTG}_{i,t} + \pi_2\, \text{LTG}_{i,t-12} + \pi_3\, \Delta\text{LTG}_{i,t} D_i^{\text{Tech}} + \pi_4\, \text{LTG}_{i,t-12} D_i^{\text{Tech}} + u_{i,t}^{(1)}. \tag{6}$$

5 Results

This section reports two sets of findings. First, the baseline IV that links five-year forecast errors $e^{(5)}_{i,t}$ to discounted five-year log returns $r^{(5)}_{i,t}$ is estimated with firm and year fixed effects (construction and identification are in Sections 3–4). Second, the possibility that the coefficients differ between technology and non-technology firms is examined via a pooled interaction IV. Before turning to the tables, Figure 2 summarizes binscatter robustness check that motivate the IV design: the left plot shows that higher $\mathrm{LTG}_{i,t-12}$ associates with more positive $e^{(5)}$; the middle plot shows a similar relation for ΔLTG; and the right plot shows a positive association between $e^{(5)}$ and subsequent $r^{(5)}$, consistent with a belief-correction channel.

Table 3 summarizes the core estimates. Columns (1)–(2) present the first-stage projections and an auxiliary reduced-form check; column (3) reports the 2SLS estimate of the effect of $e^{(5)}_{i,t}$ on $r^{(5)}_{i,t}$. All columns include firm and year fixed effects.

The first stage is strong: both ΔLTG and $\mathrm{LTG}_{i,t-12}$ predict the forecast error (col. 1), and the 2SLS column shows a Kleibergen–Paap rk Wald F of 71.53. The IV slope on $e^{(5)}_{i,t}$ is positive and statistically significant (0.2978^b, s.e. 0.1262), implying that when long-run beliefs are overly optimistic, subsequent five-year returns are higher as belief errors dissipate.

The next step is to test whether the return–forecast-error slope differs by technology status. The specification instruments both $e^{(5)}_{i,t}$ and $e^{(5)}_{i,t} \times D^{\mathrm{Tech}}_i$ with the LTG instruments and their interactions (Section 4). Table 4 reports four standard fixed-effects configurations.

In the interaction IV, the non-technology slope (the coefficient on $e^{(5)}$) is positive and conventionally significant once firm fixed effects are included; the technology differential (the coefficient on $e^{(5)} \times$ Tech) is positive but not statistically significant. System-wide first-stage strength is acceptable; in the specification with both firm and year fixed effects, the Kleibergen–Paap rk Wald F is 16.27. The Wald test yields $p = 0.1041$, which is borderline at the 10% level, so there is no statistically significant evidence that the $e^{(5)}$–$r^{(5)}$ relationship differs between technology and non-technology firms in the matched sample.

6 Robustness Checks

This section reports robustness analyses conducted to assess the stability of the baseline findings under reasonable variations of horizon, sample window, sector comparability, and instrument strength. Across all exercises, the timing convention (three-month EPS information lag), month-wise 1%/99% winsorization, fixed-effects structure, and Driscoll–Kraay standard errors are held constant, so that contrasts isolate the intended dimension of variation.

A horizon-sensitivity assessment examines whether the long-horizon mechanism is specific to five years or also appears at a shorter horizon. Under the belief-correction interpretation, the coefficient on the instrumented forecast error is expected to retain a positive sign across horizons, with statistical significance indicating robustness. The exercise reconstructs discounted three-year log returns $r_{i,t}^{(3)}$ and the corresponding three-year forecast error $e_{i,t}^{(3)}$ analogously to $r_{i,t}^{(5)}$ and $e_{i,t}^{(5)}$, applies year fixed effects, and uses Driscoll–Kraay standard errors with a 36-month bandwidth (Table 6). The resulting IV coefficient is positive and statistically significant, while the technology interaction remains not statistically significant across fixed-effects configurations (Table 7). The conclusion is that the central relation between long-horizon forecast errors and subsequent returns is not horizon-specific, and the absence of a reliable technology differential is consistent across horizons.

A subperiod-exclusion robustness check evaluates sensitivity to well-known high-volatility windows associated with the late-1990s technology boom and the 2007-2009 Global Financial Crisis. If the baseline result were driven predominantly by those episodes, the coefficient would materially weaken or change sign when such windows are removed. The specification re-estimates the five-year IV with firm and year fixed effects after excluding 1998–2002 and 2007–2009 (Table 5). The estimated coefficient remains positive and statistically significant, with an attenuated magnitude relative to the full sample, and first-stage strength remains high. The conclusion is that the principal finding is not an artifact of those subperiods.

A covariate-match assessment addresses sector comparability in the heterogeneity analysis. The rationale is that technology and non-technology firms differ systematically in observables (size, valuation, returns, leverage); propensity-score matching renders the groups more comparable, so the interaction IV is interpretable as a sectoral differential rather than a composition effect. The matching uses 1:1 nearest neighbor on firm-level time-averaged size (market capitalization), valuation (book-to-price), average monthly return, and leverage, and the interaction IV is re-estimated on the matched universe (Tables 2 and 4). After matching (for example, the leverage gap shrinks from roughly 51% to 4.5%), the non-technology slope on $e_{i,t}^{(5)}$ remains positive and statistically significant once firm fixed effects are included, and the differential $e_{i,t}^{(5)} \times D_i^{\text{Tech}}$ is positive but not statistically significant. The conclusion is that, within a matched comparison, the main effect persists while the technology differential is not statistically distinguishable from zero.

An instrument-strength and direction robustness check verifies that identification does not rely on weak instruments and that reduced-form associations align with the structural interpretation. Under standard criteria, sufficiently large first-stage statistics indicate adequate relevance, and reduced-form signs consistent with the structural estimates reinforce credibility. The baseline five-year IV with firm and year fixed effects exhibits a strong first stage (Kleibergen–Paap rk Wald $F = 71.53$; Table 3), and the pooled interaction IV retains acceptable strength (KP rk $F = 16.27$; Table 4). The auxiliary projection of $r_{i,t}^{(5)}$ on the instruments displays reduced-form directions consistent with the IV estimates (Table 3, column 2). The conclusion is that instrument relevance is strong and the sign patterns are coherent with the proposed mechanism.

Taken together, the horizon, subperiod-exclusion, matched-sample, and instrument-strength robustness checks support the stability of the core conclusion: instrumented long-horizon forecast errors positively predict subsequent discounted long-horizon returns, and the technology–non-technology slope differential is not statistically distinct within the matched universe.

7 Conclusion

This essay examines a central question with long-horizon implications: when analysts' long-term growth beliefs overshoot reality, do prices subsequently retrace, and is that corrective pattern stronger in technology? The evidence supports the first and is inconclusive on the second. Using a five-year horizon, a disciplined timing convention for building forecast errors, and LTG-based instruments within a firm–year fixed-effects panel, the baseline IV reveals a *positive, statistically significant* slope of long-run returns on long-run forecast errors. First-stage robustness checks are strong, and the effect remains positive, though attenuated, when excluding classic bubble/stress windows.

The technology–non-technology comparison is more nuanced. Allowing the slope on forecast errors to differ by technology status and instrumenting both the level and the interaction yields acceptable first-stage strength. In that overidentified system, conventional point estimates suggest a steeper relationship within technology, but standard errors are wide and the Wald test does not reject equal slopes ($p = 0.10$). The matched-sample design ensures that technology and non-technology firms are comparable on long-run observables, and fixed effects purge time-invariant heterogeneity and common shocks; even so, the data are insufficient to estimate a statistically significant differential at conventional levels.

Two limitations are noteworthy. First, measurement is demanding at long horizons: requiring strictly positive EPS at endpoints, using a three-month information lag, and winsorizing within months improves comparability but trims the sample and may attenuate estimates. Second, IV validity rests on exclusion after firm and time effects; interaction specifications also entail weaker first-stage strength. These cautions motivate extensions rather than undermine the core message. Natural next steps are to incorporate additional belief proxies (dispersion, recommendations) into a joint system; to study state dependence by interacting instruments with aggregate sentiment or funding conditions; to replicate across geographies; and to connect return predictability more tightly to subsequent earnings news versus discount-rate components. Each of these would sharpen where and when belief overreaction is most consequential.

In sum, the long-run link from forecast errors to returns survives a stringent panel-IV design and standard robustness. Whether technology systematically amplifies that link remains an open question in need of sharper instruments or richer settings. An implication for practice is that monitoring the long-horizon narrative embedded in LTG and its revisions may be informative, not because it is always wrong, but because when it errs, the subsequent adjustment can be slow, material, and—crucially—predictable.

References

Baker, Malcolm and Jeffrey Wurgler (2006). Investor sentiment and the cross-section of stock returns. In: *Journal of finance* 61.4, pp. 1645–1680. DOI: 10.1111/j.1540-6261.2006.00885.x.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer (2019). Diagnostic expectations and stock returns. In: *Journal of finance* 74.6, pp. 2839–2874. DOI: 10.1111/jofi.12833.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer (2022). Belief overreaction and stock market puzzles. In: *Journal of political economy* 132.5, pp. 1450–1484. DOI: 10.1086/727713.

Bradshaw, Mark T. (2004). How do analysts use their earnings forecasts in generating stock recommendations? In: *The accounting review* 79.1, pp. 25–50. DOI: 10.2308/accr.2004.79.1.25.

Chan, Louis K. C., Jason Karceski, and Josef Lakonishok (2003). The level and persistence of growth rates. In: *Journal of finance* 58.2, pp. 643–684. DOI: 10.1111/1540-6261.00540.

Diether, Karl B., Christopher J. Malloy, and Anna Scherbina (2002). Differences of opinion and the cross section of stock returns. In: *Journal of finance* 57.5, pp. 2113–2141. DOI: 10.1111/0022-1082.00490.

Driscoll, John C. and Aart C. Kraay (1998). Consistent covariance matrix estimation with spatially dependent panel data. In: *Review of economics and statistics* 80.4, pp. 549–560.

Frisch, Ragnar and Frederick V. Waugh (1933). Partial time regressions as compared with individual trends. In: *Econometrica* 1.4, pp. 387–401.

Ho, Daniel E., Kosuke Imai, Gary King, and Elizabeth A. Stuart (2007). Matching as nonparametric preprocessing for reducing model dependence in parametric causal inference. In: *Political analysis* 15.3, pp. 199–236.

Ho, Daniel E., Kosuke Imai, Gary King, and Elizabeth A. Stuart (2011). Matchit: nonparametric preprocessing for parametric causal inference. In: *Journal of statistical software* 42.8, pp. 1–28.

Kleibergen, Frank and Richard Paap (2006). Generalized reduced rank tests using the singular value decomposition. In: *Journal of econometrics* 133.1, pp. 97–126.

La Porta, Rafael (1996). Expectations and the cross-section of stock returns. In: *Journal of finance* 51.5, pp. 1715–1742. DOI: 10.1111/j.1540-6261.1996.tb05223.x.

Lakonishok, Josef, Andrei Shleifer, and Robert W. Vishny (1994). Contrarian investment, extrapolation, and risk. In: *Journal of finance* 49.5, pp. 1541–1578. DOI: 10.1111/j.1540-6261.1994.tb04772.x.

Lamont, Owen A. and Richard H. Thaler (2003). Can the market add and subtract? mispricing in tech-stock carve-outs. In: *Journal of political economy* 111.2, pp. 227–268. DOI: 10.1086/367678.

Lovell, Michael C. (1963). Seasonal adjustment of economic time series and multiple regression analysis. In: *Journal of the american statistical association* 58.304, pp. 993–1010.

Rosenbaum, Paul R. and Donald B. Rubin (1983). The central role of the propensity score in observational studies for causal effects. In: *Biometrika* 70.1, pp. 41–55.

Shleifer, Andrei and Robert W. Vishny (1997). The limits of arbitrage. In: *Journal of finance* 52.1, pp. 35–55. DOI: 10.1111/j.1540-6261.1997.tb03807.x.

Stambaugh, Robert F., Jianfeng Yu, and Yu Yuan (2012). The short of it: investor sentiment and anomalies. In: *Journal of financial economics* 104.2, pp. 288–302. DOI: 10.1016/j.jfineco.2011.12.001.

Wooldridge, Jeffrey M. (2010). *Econometric analysis of cross section and panel data*. 2nd ed. Cambridge, MA: MIT Press.

Tables

Table 1:
Summary statistics after matching: technology versus non-technology firms

Variable	**Group**	**Count**	**Mean**	**SD**	**P10**	**Median (P50)**	**P90**
LTG	Non-Technology	33,885	13.51	8.07	6.50	12.40	21.67
	Technology	29,862	17.68	9.37	9.00	16.17	27.99
$e^{(5)}$	Non-Technology	33,885	-5.33	15.59	-22.43	-3.74	9.44
	Technology	29,862	-7.10	19.96	-30.86	-5.19	12.57
$e^{(3)}$	Non-Technology	33,049	-4.21	21.39	-25.89	-2.55	14.98
	Technology	27,982	-5.23	28.06	-35.60	-2.88	20.36
$r^{(3)}$	Non-Technology	33,885	28.13	51.46	-29.46	30.31	86.59
	Technology	29,862	28.23	64.65	-46.64	29.60	103.35
$r^{(5)}$	Non-Technology	33,885	43.71	63.79	-27.66	46.14	115.30
	Technology	29,862	45.41	76.56	-47.82	46.08	138.37
Return	Non-Technology	32,630	1.31	9.93	-9.53	1.24	11.95
	Technology	29,100	1.53	13.23	-13.32	1.23	16.30
EPS (LTM)	Non-Technology	33,885	2.43	8.35	0.19	1.20	4.27
	Technology	29,862	1.49	2.29	0.09	0.81	3.52

Notes. This table reports distributional summaries for the matched technology and non-technology samples (Section 3). Statistics are computed cross-sectionally across firm–months *after* matching. Entries for LTG, Error, Error3, $r^{(3)}$, $r^{(5)}$, and Return are shown in *percent* (i.e., original decimals multiplied by 100); interpret these as *percentage points*. $r^{(3)}$ and $r^{(5)}$ are discounted log returns constructed from forward monthly returns in annual blocks; because they are log aggregates, percent values can exceed 100. EPS (LTM) is trailing-twelve-month EPS (I/B/E/S actuals) in USD per share and is not converted to percent. Differences in counts between technology and non-technology primarily reflect heterogeneous time coverage across firms (entry/exit and data availability). All variables are winsorized at the 1st/99th percentiles by calendar month prior to aggregation.

Table 2:
Comparison of means and sample sizes: technology versus non-technology firms (before and after matching)

Variable	Stage	Count (Non-Technology)	Mean (Non-Technology)	Count (Technology)	Mean (Technology)	% Diff vs. Technology
Size (Market Cap)	Before	266,359	10.20	37,740	12.80	-20.31%
	After	33,626	11.90	35,150	17.00	-30.00%
Valuation (B/P)	Before	293,327	0.585	40,513	0.495	+18.18%
	After	37,984	0.562	38,697	0.487	+15.26%
Return (%)	Before	299,643	1.14	41,106	1.44	-20.29%
	After	38,481	1.24	39,076	1.52	-18.04%
Leverage (D/A)	Before	299,617	0.262	36,306	0.173	+51.05%
	After	36,252	0.181	33,466	0.173	+4.54%

Notes. This table reports mean characteristics and counts for technology and non-technology firms before and after propensity-score matching. Counts are firm–month observations; sample sizes differ by row because of variable-specific missingness. Size is the time-averaged market capitalization in **billions of USD**. Valuation is the time-averaged book-to-price ratio (book value per share divided by price). Return is the time-averaged monthly return shown in **percentage points**. Leverage is the time-averaged total debt-to-total assets ratio (D/A). All variables are winsorized at the 1% and 99% levels before aggregation.

Percent difference is calculated as

$$\% \text{ Diff} = \frac{\text{Non-Technology Mean} - \text{Technology Mean}}{\text{Technology Mean}} \times 100\,.$$

Table 3:
Five-year horizon IV with firm and year fixed effects

VARIABLES	(1) First stage: $e^{(5)}$	(2) Aux. check ($r^{(5)}$ on IVs)	(3) 2SLS: $r^{(5)}$
ΔLTG	-0.2840^{a}	-0.0817^{b}	
	(0.0239)	(0.0367)	
$\text{LTG}_{i,t-12}$	-0.2779^{a}	-0.0859^{c}	
	(0.0393)	(0.0462)	
$e^{(5)}$			0.2978^{b}
			(0.1262)
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
N	52,734	52,734	52,734
KP rk F			71.5310

Notes: This table reports the baseline instrumental-variables (IV) estimates linking five-year forecast errors to discounted five-year log returns at the firm level. Standard errors are reported in brackets. Outcome $r^{(5)}$ is the discounted five-year log return; endogenous regressor $e^{(5)}$ is the five-year EPS forecast error (Sections 3–4). Excluded instruments are ΔLTG and $\text{LTG}_{i,t-12}$. All specifications include firm and year fixed effects; standard errors are Driscoll–Kraay (DK) with long bandwidth (60 months). "Aux. check" in col. (2) projects $r^{(5)}$ on the instruments to gauge reduced-form direction; it is not used for inference. Significance: $^{a}p<0.01$, $^{b}p<0.05$, $^{c}p<0.10$. †As a rule of thumb, KP rk $F > 10$ indicates strong first-stage relevance.

Table 4:
Interaction IV — comparison of fixed-effects specifications

VARIABLES	(1) $r^{(5)}$	(2) $r^{(5)}$	(3) $r^{(5)}$	(4) $r^{(5)}$
$e^{(5)}$	0.0308	−0.0784	0.1487^{c}	0.1545^{b}
	(0.0806)	(0.0817)	(0.0875)	(0.0767)
$e^{(5)} \times$ Tech	−0.0872	−0.1036	0.2954	0.3364
	(0.2577)	(0.2864)	(0.2190)	(0.2070)
Firm FE	No	No	Yes	Yes
Year FE	No	Yes	No	Yes
N	52,739	52,739	52,739	52,739
KP rk F	15.1103	18.6049	11.4532	16.2749

Notes: This table reports pooled interaction IV estimates that allow the slope on the forecast-error term to differ between technology and non-technology firms under alternative fixed-effects configurations. Standard errors are reported in brackets. Outcome $r^{(5)}$. Endogenous regressors are $e^{(5)}$ and $e^{(5)} \times D_i^{\text{Tech}}$. Excluded instruments: $\{\Delta$LTG, LTG$_{i,t-12}$, ΔLTG$\times$Tech, LTG$_{i,t-12}\times$Tech$\}$. FE rows indicate the fixed effects in each column. Standard errors are Driscoll–Kraay (DK) (bandwidth 60). The matched-sample universe is used for the technology versus non-technology comparison. Significance: $^{a}p<0.01$, $^{b}p<0.05$, $^{c}p<0.10$. †As a rule of thumb, KP rk $F > 10$ indicates strong first-stage relevance.

Table 5:
Baseline five-year horizon IV (full sample versus excluding 1998–2002 and 2007–2009)

VARIABLES	(1) First stage: $e^{(5)}$	(2) Aux. check ($r^{(5)}$ on IVs)	(3) 2SLS: $r^{(5)}$	(4) 2SLS: $r^{(5)}$ (excl. windows)
ΔLTG	-0.2840^{a}	-0.0817^{b}		
	(0.0239)	(0.0367)		
$\text{LTG}_{i,t-12}$	-0.2779^{a}	-0.0859^{c}		
	(0.0393)	(0.0462)		
$e^{(5)}$			0.2978^{b}	0.1443^{c}
			(0.1262)	(0.0772)
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
N	52,734	52,734	52,734	38,254
KP rk F			71.5310	86.5301

Notes: This table reports baseline five-year IV results for the full sample and for a subsample that excludes the 1998–2002 and 2007–2009 windows of extreme volatility. Standard errors are reported in brackets. Outcome $r^{(5)}$. Endogenous regressor $e^{(5)}$. Excluded instruments: ΔLTG, $\text{LTG}_{i,t-12}$. FE and standard errors as in the main text. Significance: $^{a}p<0.01$, $^{b}p<0.05$, $^{c}p<0.10$. †As a rule of thumb, KP rk $F > 10$ indicates strong first-stage relevance.

Table 6:
Baseline three-year horizon IV (full sample versus excluding 1998–2002 and 2007–2009)

VARIABLES	(1) First stage: $e^{(3)}$	(2) Aux. check ($r^{(3)}$ on IVs)	(3) 2SLS: $r^{(3)}$	(4) 2SLS: $r^{(3)}$ (excl. windows)
ΔLTG	-0.1097^a	-0.0970^a		
	(0.0275)	(0.0340)		
LTG$_{i,t-12}$	-0.1116^a	-0.0914^b		
	(0.0363)	(0.0431)		
$e^{(3)}$			0.8517^a	0.4590^b
			(0.3049)	(0.2242)
Firm FE	No	No	No	No
Year FE	Yes	Yes	Yes	Yes
N	53,176	53,176	53,176	39,018
KP rk F			8.4119	5.3178

Notes: This table reports instrumental-variables estimates using three-year forecast errors and discounted three-year log returns, analogous to the five-year baseline specification. Standard errors are reported in brackets. Outcome $r^{(3)}$ is the discounted three-year log return; endogenous regressor $e^{(3)}$ is the three-year EPS forecast error constructed analogously to $e^{(5)}$ but over a three-year horizon. Excluded instruments are $\Delta\text{LTG} \equiv \text{LTG}_t - \text{LTG}_{t-12}$ and $\text{LTG}_{i,t-12}$. All specifications include *year* fixed effects only. Standard errors are Driscoll–Kraay with a 36-month bandwidth. Column (4) excludes 1998–2002 and 2007–2009. Preprocessing (winsorization/standardization) follows the main text. Significance: $^a p<0.01$, $^b p<0.05$, $^c p<0.10$. †As a rule of thumb, KP rk $F > 10$ indicates strong first-stage relevance.

Table 7:
Interaction IV — comparison of fixed-effects specifications (three-year horizon)

VARIABLES	(1) $r^{(3)}$	(2) $r^{(3)}$	(3) $r^{(3)}$	(4) $r^{(3)}$
$e^{(3)}$	0.1126	−0.0854	0.5012^{a}	0.6594^{a}
	(0.2208)	(0.2096)	(0.1589)	(0.1952)
$e^{(3)} \times$ Tech	−0.1381	−0.1783	0.0935	0.3907
	(0.4273)	(0.4680)	(0.7223)	(0.7363)
Firm FE	No	No	Yes	Yes
Year FE	No	Yes	No	Yes
N	53,179	53,179	53,179	53,179
KP rk F	4.0312	3.5119	0.6423	1.1539

Notes: This table reports pooled interaction IV estimates at the three-year horizon, comparing technology and non-technology slopes under alternative fixed-effects structures. Standard errors are reported in brackets. Outcome $r^{(3)}$. Endogenous regressors are $e^{(3)}$ and $e^{(3)} \times D_i^{\text{Tech}}$. Excluded instruments: $\{\Delta\text{LTG}, \text{LTG}_{i,t-12}, \Delta\text{LTG} \times D_i^{\text{Tech}}, \text{LTG}_{i,t-12} \times D_i^{\text{Tech}}\}$. Fixed-effects configurations are indicated in the two FE rows. Standard errors are Driscoll–Kraay with a 36-month bandwidth. The matched-sample universe is used for the technology versus non-technology comparison. Significance: $^{a}p<0.01$, $^{b}p<0.05$, $^{c}p<0.10$. †As a rule of thumb, KP rk $F > 10$ indicates strong first-stage relevance.

Figures

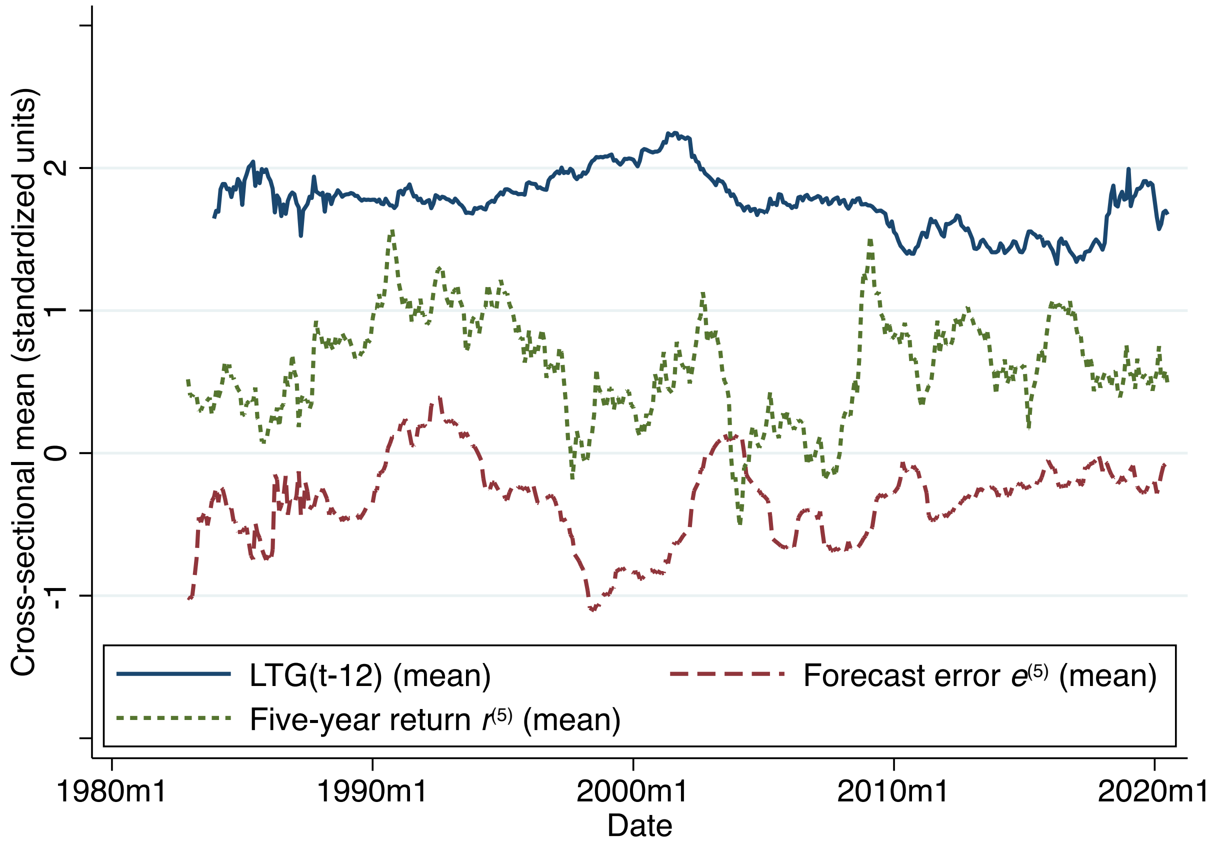

Figure 1:
Time-series of cross-sectional means: $\text{LTG}_{i,t-12}$, $e^{(5)}$, and $r^{(5)}$.
Notes: This figure tracks the evolution of average long-term growth forecasts ($\text{LTG}_{i,t-12}$), forecast errors ($e^{(5)}$), and discounted five-year returns ($r^{(5)}$) across months. The comovement between $e^{(5)}$ and $r^{(5)}$ illustrates the inverse relation between optimism and subsequent performance. Series are monthly cross-sectional means computed from the matched-sample universe. Variables follow the paper's preprocessing: month-wise winsorization at the 1st/99th percentiles and, where indicated, standardization by the full-sample standard deviation. Vertical axis units are standardized deviations from the full-sample mean.

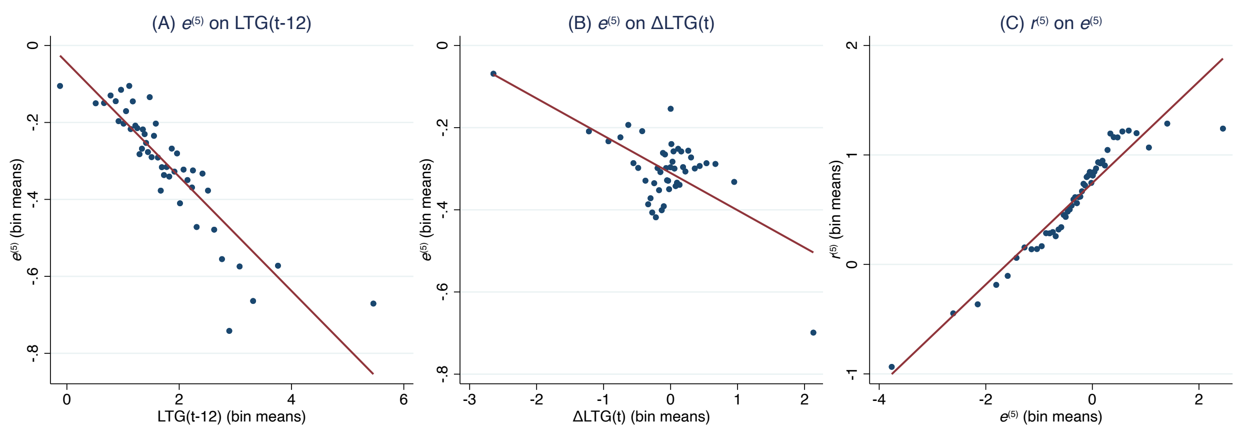

Figure 2:
Binscatter for the IV model: (i) $e^{(5)}$ on $\text{LTG}_{i,t-12}$, (ii) $e^{(5)}$ on ΔLTG, and (iii) $r^{(5)}$ on $e^{(5)}$.
Notes: Each panel plots a binned scatter graph, where the sample is divided into equal-frequency bins along the x-axis, and the plotted points represent mean x–y values within each bin. The fitted line represents the ordinary least squares (OLS) regression on the underlying data, not on the binned means. Panels (A) and (B) show the first-stage relationships used as instruments, both indicating strong negative correlations between long-term growth expectations and subsequent forecast errors. Panel (C) displays the reduced-form link between forecast errors and long-horizon returns, highlighting the negative association central to the IV design. Variables are winsorized by month at the 1st/99th percentiles and standardized by the full-sample standard deviation.

Appendix

Table 8:
Core variables, definitions, and sources

Internal name	Definition (operational)	Source / Units
`LTG`	Analysts' long-term EPS growth at t; percent converted to decimal.	LSEG `TR.LTGMean` / decimal
`EPS (LTM)`	Trailing-twelve-month $\text{EPS}_{i,t}$.	LSEG `TR.EPSActValue` / native
`Total return`	Simple monthly return $r_{i,t}$; percent converted to decimal.	LSEG `TR.TotalReturn` / decimal
`Five-year return`	Discounted five-year log return ($r^{(5)}$).	Derived / decimal
`Five-year error`	Five-year error $e^{(5)}$ is the difference between the LTG and the actual EPS in that time horizon: $[\ln(\text{EPS}_{t+57}) - \ln(\text{EPS}_{t-3})]/5 - \text{LTG}_t$.	Derived / decimal
`Price`	Month-end closing price.	LSEG `TR.PriceClose` / native
`BVPS`	Book value per share.	LSEG `TR.F.BookValuePerShr` / native
`MktCap`	Company market capitalization.	LSEG `TR.CompanyMarketCap` / USD
`GICS group`	Industry group code (categorical).	LSEG `TR.GICSIndustryGroupCode` / code
`Technology dummy`	$D_i^{\text{Tech}} = 1$ if GICS group is one of Software & Services (4510), Technology Hardware & Equipment (4520), or Semiconductors & Semiconductor Equipment (4530), else 0.	Derived / binary

Notes: This table reports the definitions and data sources for core variables used in the analysis. All logs are natural logs. Winsorization is applied within calendar month at the 1st/99th percentiles to LTG, forecast errors, long-horizon returns, and monthly return. When variables are standardized for interpretation, instruments are formed from the standardized LTG series.