Cash-Flow Duration and the Interest-Rate Sensitivity of Anomaly Returns

[Author]*

April 9, 2026

Abstract

A single characteristic forecasts the cross-section of anomaly interest-rate betas: the long-short book-to-market spread of the portfolio's legs. Across 135 Chen–Zimmerman anomaly signals, the cross-sectional regression slope is +1.43 with heteroskedasticity-robust $t = +5.25$ ($R^2 = 0.14$); on the 124-signal subsample that excludes explicit value signals the slope is +2.31 with $t = +4.93$. The pattern survives orthogonalized federal funds rate shocks, forward-looking one-year treasury shocks, joint controls for credit, VIX, and TIPS-breakeven exposure, and three of four ex-value definitions. A two-stage cash-flow duration model supplies an interpretive scaffolding: rate-beta equals the leg-weighted modified-duration difference, and under Gordon growth with bivariate normality the cross-sectional regression slope is a linear projection. The theoretical scaling constant is roughly an order of magnitude above the empirical slope, and a calibration exercise shows multi-stage duration cannot close the gap at any cross-firm terminal-growth dispersion consistent with IBES long-term-growth forecasts. The residual magnitude gap is the paper's open puzzle. Book-to-market spread explains 14% of the cross-sectional variation in rate-beta; the remaining 86% belongs to signal-specific factors the paper does not model.

JEL codes: G12, G14, E43.

Keywords: cash-flow duration, interest rates, anomalies, book-to-market.

*[Affiliation]. Email: [email]. I thank [acknowledgements].

1 Introduction

From January 2022 to December 2024 the federal funds rate climbed roughly 525 basis points, ending a fifteen-year near-zero regime. The Chen–Zimmerman anomaly cross-section, dormant since the early 2010s, responded sharply. Some signals revived. Others did not. A natural question follows. Which anomalies should respond to interest rates, and does a simple observable characteristic forecast the response?

The long-short book-to-market (BM) spread of an anomaly portfolio's legs predicts that portfolio's interest-rate beta across the Chen–Zimmerman universe. The cross-sectional regression of signal rate-beta on BM-spread delivers a slope of +1.43 with heteroskedasticity-robust $t = +5.25$ on 135 signals ($R^2 = 0.14$), and +2.31 with $t = +4.93$ on the 124-signal subsample that excludes the explicit value family. The pattern survives every identification stress applied: orthogonalized federal funds rate shocks, forward-looking one-year treasury shocks, joint cross-sectional controls for credit, VIX, and TIPS-breakeven exposure, and three of four ex-value sample definitions.

The empirical mechanism operates through duration mismatch at the portfolio level. A long-short anomaly inherits a duration mismatch from its leg composition: any signal whose long leg holds shorter-duration stocks than its short leg picks up positive rate beta. High-BM stocks are short-duration in the Gordon limit because $1/(r - g)$ falls as the implied growth rate g falls. The BM spread of a signal's legs measures composition duration directly, and the data confirm the sign and the ordering.

Scope of the empirical claim. The BM-spread regressor accounts for 14% of the cross-sectional variation in rate-beta. The remaining 86% is signal-specific variance the paper does not model. A single cross-sectional characteristic captures a material fraction of the anomaly rate-beta cross-section, and the fraction it captures is orthogonal to credit, VIX, and inflation-expectation exposures. BM-spread does not span the cross-section of rate betas; it predicts a non-trivial component of it with $t = +5.25$.

Theory as interpretive scaffolding. Section 2 develops a two-stage cash-flow duration framework as an interpretive scaffolding. The framework does three things. First, Proposition 1 writes signal rate-beta as the leg-weighted modified-duration difference, a standard accounting identity under a single-factor rate shock. Second, Proposition 2 writes the Gordon rate-beta as a linear projection on BM-spread, which fixes the sign of the cross-sectional regression slope but not its quantitative magnitude. Third, Calibration Result 1 (a finite-difference calculation at three points, not an analytic characterization) asks whether multi-stage duration can close the order-of-magnitude gap between the theoretical Gordon scaling constant and the empirical slope. It cannot.

The quantitative gap is the paper's open puzzle. Gordon cash flows at $r = 0.04$, $\bar{g} = 0.01$ deliver a theoretical slope of $K_G = 33.3$; the empirical slope is $\hat{K}_{\text{emp}} = 0.94$; the ratio is $K_G/\hat{K}_{\text{emp}} = 35.4$. The natural multi-stage refinement (Weber two-stage with finite growth phase and terminal perpetuity) compresses the slope only when the cross-firm dispersion of terminal growth Δg_L is essentially zero, and direct measurement of analyst long-term growth dispersion on 842,290 IBES firm-months shows $\sigma(g_L) \approx 8.94$ percentage points, far from zero. Multi-stage duration cannot close the gap.

Interpretation of the magnitude gap. One plausible reading has two components. Proposition 2 is a first-order linear projection of rate-beta on the BM-spread regressor. The empirical slope $\hat{K}_{\text{emp}}$ captures the component of rate-beta that loads linearly on BM-spread, so the $K_G/\hat{K}_{\text{emp}}$ ratio reflects (i) classical measurement-error attenuation because book-to-market is a noisy proxy for true cash-flow duration and (ii) higher-order terms orthogonal to BM-spread that the first-order projection does not capture. A quantitative anchor is available for the first channel only. **?** reports firm-level correlations between book-to-market and IBES-based duration in the $\rho \in [0.3, 0.5]$ range; under classical errors-in-variables the attenuation factor on the regression slope is $1/\rho^2 \in [4, 11]$. Measurement error on this scale accounts for roughly one-third of the $35\times$ ratio at the upper end ($\rho = 0.3$, giving $11/35 \approx 32\%$) and roughly a ninth at the lower end ($\rho = 0.5$, giving $4/35 \approx 11\%$). The residual unexplained factor is roughly 3 to $9\times$ and must come from higher-order terms orthogonal to BM-spread. The paper does not decompose that residual further and does not claim the Gordon benchmark K_G is an unbiased point estimate of the duration channel. Proposition 3 below is a robustness exercise for this reading: the Edgeworth bound shows the bivariate-normality assumption in Proposition 2 is not load-bearing for the functional form, so the linear projection survives non-normal sorting characteristics with bounded kurtosis.

Relation to prior work. **?** construct duration-matched bond benchmarks for 246 anomaly portfolios and evaluate whether an anomaly's expected return survives net of its duration hedge. Their object is expected returns, and the duration channel enters as a nuisance to be hedged out: the goal is to neutralize it. This paper's object is different. The focus is the cross-section of rate-betas themselves, and the BM-spread characteristic serves as a predictor rather than a hedge instrument. Both papers rely on the same underlying accounting fact (that equity duration is heterogeneous and BM-correlated), but the two research programs are complementary: van Binsbergen et al. seek to remove the duration channel from anomaly returns via benchmarks, while this paper exposes and measures the channel as a cross-sectional forecaster. The van Binsbergen et al. supplementary analysis does not report a cross-sectional regression of anomaly rate-beta on BM-spread in the form of equation (21); nor does it confront the magnitude gap between the theoretical Gordon scaling constant and the empirical slope. **?** estimates firm-level multi-stage durations from IBES long-term growth forecasts and documents firm-level rate sensitivities; this paper uses the same IBES series differently, as an input to a calibration sanity check at the portfolio level. **?** use dividend-futures data to explain expected return levels of five named factors; this paper's object is rate-beta across 135 Chen–Zimmerman signals, not expected returns across a handful of factors. **?** documents a short-duration premium at the firm level; this paper remains agnostic on the sign of expected return premia and documents a predictability fact about rate-beta at the portfolio level. **?** develop the theoretical benchmark for duration-based pricing of the value premium; this paper's projection identity shares the duration intuition but operates at the portfolio level across 135 signals rather than at the level of the value premium itself.

Roadmap. Section 2 sets up the duration-mismatch framework, derives the projection identity, and reports the multi-stage calibration sanity check. Section 3 presents the main cross-sectional regression, the identification robustness, the ex-value sensitivity, the IBES long-term-growth moments, and an out-of-sample check on the 2022–2024 hiking cycle. Section 4 discusses implications and limitations. Section 5 concludes. Appendix A collects the algebraic details and calibration arithmetic.

2 Model

The duration-mismatch framework in this section supplies an interpretive scaffolding for the empirical results in Section 3. Subsection 2.2 derives the rate-beta/leg-duration identity. Subsection 2.3 writes the Gordon rate-beta as a first-order linear projection on the BM-spread. Subsection 2.4 bounds the projection error under a weaker moment assumption. Subsection 2.5 gives an existence argument for why the projection identity cannot span the full cross-section. Subsection 2.6 is a calibration sanity check: multi-stage duration does not close the gap between the theoretical Gordon scaling constant and the empirical slope. Subsection 2.7 measures the relevant IBES moments and confirms that the sanity check fails as expected.

None of the results in this section is new. Each formalizes an accounting identity, a textbook duration formula, or a back-of-envelope calibration.

2.1 Setup

Time is discrete, $t = 0, 1, 2, \ldots$. A cross-section of N firms indexed by i holds cash-flow streams $\{CF_{i,t}\}$. The discount factor is $1 + r$ with $r > 0$ common across firms. Firm value satisfies

$$V_i(r) = \sum_{t=1}^{\infty} \frac{CF_{i,t}}{(1+r)^t}. \tag{1}$$

Two specializations follow.

Gordon growth. Cash flows follow a constant growth path, $CF_{i,t} = CF_{i,0}(1 + g_i)^t$ with $g_i < r$. Equation (1) simplifies to

$$V_i(r) = \frac{CF_{i,0}(1 + g_i)}{r - g_i}, \tag{2}$$

so that

$$\text{BM}_i \equiv \frac{B_i}{V_i} = \frac{B_i(r - g_i)}{CF_{i,0}(1 + g_i)}. \tag{3}$$

Weber two-stage. For $t \le T_i$, $CF_{i,t} = CF_{i,0}(1 + g_{H,i})^t$. For $t > T_i$, $CF_{i,t} = CF_{i,T_i}(1 + g_{L,i})^{t-T_i}$ with $g_{L,i} < r$. Firm value is

$$V_i(r) = CF_{i,0}\left[\frac{(1 + g_{H,i})(1 - \rho_{H,i}^{T_i})}{r - g_{H,i}} + \rho_{H,i}^{T_i} \cdot \frac{1 + g_{L,i}}{r - g_{L,i}}\right], \quad \rho_{H,i} \equiv \frac{1 + g_{H,i}}{1 + r}. \tag{4}$$

The Gordon specification (2) emerges as $T_i \to \infty$ with $g_i = g_{H,i}$, and a pure mature perpetuity at $g_{L,i}$ emerges as $T_i \to 0$.

Rate shocks. Let Δr be a scalar perturbation to r realized one period after valuation. Firm i's one-period return is

$$R_i = \frac{V_i(r + \Delta r) + CF_{i,1} - V_i(r)}{V_i(r)} + \epsilon_i, \tag{5}$$

where ϵ_i is idiosyncratic cash-flow news. Maintain throughout:

(A1) Deterministic cash flows: $\text{Cov}(\epsilon_i, \Delta r) = 0$.

(A2) Single-factor rate shock: Δr is scalar.

(A3) $|\Delta r|$ is small; the first-order Taylor approximation is accurate up to $O((\Delta r)^2)$.

(A4) Cash flows do not respond to Δr.

Modified durations. The modified duration of firm i is $D_i^{\mathrm{mod}} \equiv -(1/V_i)\, dV_i/dr$. Under Gordon,

$$D_{i,G}^{\mathrm{mod}} = \frac{1}{r - g_i}, \tag{6}$$

derived in Appendix A.1. Under Weber two-stage,

$$D_{i,W}^{\mathrm{mod}} = w_{1,i} D_{1,i}^{\mathrm{mod}} + w_{2,i} D_{2,i}^{\mathrm{mod}}, \tag{7}$$

where $w_{1,i} = V_{1,i}/V_i$ and $w_{2,i} = V_{2,i}/V_i$ are the value weights of the two phases.

Lemma 1 (Phase-2 Weber modified duration)**.** *The phase-2 modified duration is*

$$D_{2,i}^{mod} = \frac{T_i}{1+r} + \frac{1}{r - g_{L,i}}, \tag{8}$$

and $0 < D_{i,W}^{mod} \leq T_i/(1+r) + 1/(r - g_{L,i})$*, with equality when* $w_{1,i} = 0$*.*

The proof is in Appendix A.2. The economic content of Lemma 1 is that $D_{2,i}^{\mathrm{mod}}$ does not depend on the phase-1 growth rate $g_{H,i}$. When firms share a common terminal growth rate g_L, the perpetuity-tail contribution to duration is identical across firms and cross-sectional duration dispersion reduces to the phase-1 annuity differential, bounded by T. Calibration Result 1 below converts this observation into a three-point numerical exercise.

Anomaly signals. An anomaly signal k is a long-short portfolio with leg weights w_k^L and w_k^S, each summing to one. The leg modified durations are

$$D_{L,k}^{\mathrm{mod}} \equiv \sum_i w_{k,i}^L D_i^{\mathrm{mod}}, \qquad D_{S,k}^{\mathrm{mod}} \equiv \sum_i w_{k,i}^S D_i^{\mathrm{mod}}. \tag{9}$$

The signal return is $R_k = \sum_i (w_{k,i}^L - w_{k,i}^S) R_i$, and its rate-beta is

$$\beta_k^r \equiv \frac{\mathrm{Cov}(R_k, \Delta r)}{\mathrm{Var}(\Delta r)}. \tag{10}$$

Log-BM decomposition. Under Gordon, taking logs in (3),

$$\log(r - g_i) = \log \mathrm{BM}_i + \xi_i, \qquad \xi_i \equiv -\log B_i + \log CF_{i,0} + \log(1 + g_i). \tag{11}$$

$\log \mathrm{BM}_i$ and ξ_i jointly determine duration; neither alone suffices. Under Weber two-stage, no such clean identity exists because V_i depends nonlinearly on $(T_i, g_{H,i}, g_{L,i})$.

2.2 Duration decomposition of rate-beta

Proposition 1 (Duration decomposition)**.** *Under (A1)–(A4), the rate-beta of any long-short signal k satisfies*

$$\beta_k^r = D_{S,k}^{mod} - D_{L,k}^{mod}. \tag{12}$$

The result holds under both the Gordon (2) *and Weber* (4) *specifications.*

Proof. Differentiating (1) in r gives $dV_i/dr = -V_i D_i^{\text{mod}}$, so by (A3),

$$\frac{V_i(r+\Delta r) - V_i(r)}{V_i(r)} = -D_i^{\text{mod}} \Delta r + O((\Delta r)^2).$$

Substitute into (5) and use (A1) to get $\text{Cov}(R_k, \Delta r) = -(D_{L,k}^{\text{mod}} - D_{S,k}^{\text{mod}}) \operatorname{Var}(\Delta r) + O(\sigma_r^3)$. Dividing by $\operatorname{Var}(\Delta r)$ yields (12). The derivation uses only (5) and $dV_i/dr = -V_i D_i^{\text{mod}}$, both of which hold for any valuation function $V_i(r)$. □

A signal whose long leg has shorter modified duration than its short leg has positive rate-beta. When rates rise, the long-duration short leg loses more, so the long-short return gains.

2.3 The projection identity under Gordon

Define the long-short decile spread of log BM as

$$\text{BM-spread}_k \equiv \mathbb{E}[\log \text{BM}_i \mid i \in L_k] - \mathbb{E}[\log \text{BM}_i \mid i \in S_k], \tag{13}$$

and likewise Δu_k for the cross-leg difference in the OLS residual from the projection

$$\xi_i = \alpha + \gamma \log \text{BM}_i + u_i, \qquad \mathbb{E}[u_i \log \text{BM}_i] = 0. \tag{14}$$

Let λ_k denote the long-short decile spread of the standardized sorting variable. Appendix A.4 shows $\lambda_k = 3.51$ for a decile sort under standard normality.

Proposition 2 (Gordon projection identity)**.** *Assume (A1)–(A4), Gordon growth* (2)*, and*

(B1) Decile sorts.

(B2) Bivariate normality: $(X_i, \log \text{BM}_i, u_i)$ *jointly normal.*

(B3) The first-order Taylor expansion of $\exp(\cdot)$ *around the cross-sectional mean of* $\log \text{BM}$ *is accurate.*

Then to leading order in $\sigma_{\log \mathrm{BM}}$*,*

$$\beta_k^r = \frac{1+\gamma}{r-\bar{g}} \cdot \textit{BM-spread}_k + \epsilon_k, \qquad \epsilon_k = \frac{\sigma_u \rho_{X,u} \lambda_k}{r-\bar{g}}, \tag{15}$$

with equality under (B1)–(B3). If X *is a deterministic linear function of* $\log \mathrm{BM}$*, then* $\rho_{X,u} = 0$ *and* $\epsilon_k = 0$*.*

Proof. From (6) and (11), $D_{i,G}^{\mathrm{mod}} = \exp(-(1+\gamma)\log \mathrm{BM}_i - \alpha - u_i)$. A first-order Taylor expansion around the cross-sectional mean of $\log \mathrm{BM}$ gives

$$D_{i,G}^{\mathrm{mod}} \approx K_G\big[1 - (1+\gamma)(\log \mathrm{BM}_i - \overline{\log \mathrm{BM}}) - (u_i - \bar{u})\big], \qquad K_G = \frac{1}{r-\bar{g}} + O(\sigma_{\log \mathrm{BM}}^2).$$

Taking the long-short decile difference and using Proposition 1,

$$\beta_k^r = D_{S,k}^{\mathrm{mod}} - D_{L,k}^{\mathrm{mod}} \approx K_G\big[(1+\gamma) \cdot \text{BM-spread}_k + \Delta u_k\big].$$

Under bivariate normality of (X, u), $\mathbb{E}[u \mid X]$ is linear, so $\mathbb{E}[u \mid i \in L_k] - \mathbb{E}[u \mid i \in S_k] = \rho_{X,u}\sigma_u\lambda_k$ exactly. □

Equation (15) decomposes rate-beta into a BM-spread-spanned component and an orthogonal residual. The scaling constant $(1+\gamma)/(r-\bar{g})$ fixes the sign of the cross-sectional regression slope. The empirical $t = +5.25$ on the full panel and $t = +4.93$ on the ex-value subsample confirm the sign.

2.4 Edgeworth bound on the projection residual

Proposition 2 is exact under bivariate normality. When (X, u) is non-normal, the conditional expectation $\mathbb{E}[u \mid X]$ need not be linear and the decile spread Δu_k picks up corrections in higher cumulants. A weaker "linear projection" assumption, combined with a Gram–Charlier approximation of the sorting density, delivers a formal bound: the skewness term cancels in the decile spread and the leading correction depends on kurtosis alone with an explicit small coefficient.

Proposition 3 (Edgeworth bound)**.** *Maintain (A1)–(A4), Gordon growth, (B1), and (B3). Replace (B2) with*

(B2′) Linear projection: $u_i = (\rho_{X,u}\sigma_u/\sigma_X)X_i + \eta_i$ *with* η_i *independent of* X_i*.*

(B2′) Gram–Charlier density of standardized X:

$$f_X(x) = \phi(x)\left[1 + \frac{\xi_3}{6}H_3(x) + \frac{\kappa_4}{24}H_4(x)\right] + o(\xi_3, \kappa_4), \tag{16}$$

where $H_3(x) = x^3 - 3x$ and $H_4(x) = x^4 - 6x^2 + 3$ are Hermite polynomials.

Then to first order in (ξ_3, κ_4),

$$\Delta u_k = \sigma_u \rho_{X,u} \lambda_k \cdot (1 + c_2\kappa_4) + O(\xi_3^2, \kappa_4^2), \tag{17}$$

where

$$c_2 = \frac{1}{24}\left[(c^4 - 2c^2 - 1) - (c^3 - 3c)\frac{\phi(c)}{1 - \Phi(c)}\right], \qquad c = \Phi^{-1}(0.9) \approx 1.2816. \tag{18}$$

Substituting numerical values yields $c_2 \approx 0.0612$, and the skewness coefficient ξ_3 drops out of the decile spread.

Appendix A.5 gives the full Hermite-integral derivation. Under linear projection, skewness in the sorting characteristic does not perturb the decile spread. Only kurtosis matters, and only at coefficient $c_2 \approx 0.0612$. For realistic excess kurtosis $\kappa_4 \in [0, 4]$, the normality error is bounded by $|c_2\kappa_4| \leq 0.244$, at most 24%. Proposition 3 is a robustness exercise for the functional form of Proposition 2. The headline regression in equation (21) does not rely on linearity for its validity as an empirical fact, and the linear projection in Proposition 2 delivers an exact equality under bivariate normality. The Edgeworth bound shows bivariate normality is not load-bearing: even if cross-sectional (X, u) is non-normal with excess kurtosis up to 4, the scaling constant in Proposition 2 is pinned down to within 24% of its normal value, and skewness drops out at first order. Non-normality of the sorting characteristic is therefore not the explanation for the 35× magnitude gap.

2.5 A constructive violator class

The projection identity in Proposition 2 must fail for sorts whose characteristic is non-monotone in $\log \text{BM}$. Proposition 4 constructs one such class explicitly.

Proposition 4 (Non-monotone violator)**.** *Under (A1)–(A4) and Gordon growth, there exists a cross-sectional joint distribution of $(\log \text{BM}, u, X)$ satisfying all assumptions of Proposition 2 such that*

(i) X is a well-defined sorting characteristic,

(ii) the decile sort on X has BM-spread$_k = 0$ exactly, and

(iii) the decile sort on X has $\beta_k^r \neq 0$ strictly.

Proof. Let $\log \mathrm{BM} \sim N(0, s^2)$ and $u \sim N(0, \sigma_u^2)$ be independent. Define

$$X_i = (\log \mathrm{BM}_i)^2 + \theta u_i, \qquad \theta > 0. \tag{19}$$

For (ii): the joint density of $(\log \mathrm{BM}, u)$ is symmetric in $\log \mathrm{BM}$. Hence the joint density of $(y, y^2 + \theta u) = (y, X)$ is symmetric in y at every fixed X. For any event $\{X \in A\}$ the density puts equal mass at (y, A) and $(-y, A)$, so $\mathbb{E}[\log \mathrm{BM} \mid X \in A] = 0$ and BM-spread$_k = 0$.

For (iii): from the Proposition 2 decomposition, $\beta_k^r = K_G \cdot \Delta u_k$. Fix $y = 0$. Conditional on $y = 0$, $X = \theta u$, so the top decile of X conditional on $y = 0$ has u bounded strictly above zero. By continuity of the joint density, a positive-measure neighborhood of $y = 0$ contributes mass to the top decile with u strictly positive on average. By symmetry, the bottom decile contributes u strictly negative. So $\Delta u_k > 0$ and $\beta_k^r > 0$. □

Section 3 reports descriptive evidence consistent with Proposition 4: momentum, volatility, and size sorts (non-monotone in $\log \mathrm{BM}$) have 4.35 times the spanning-residual variance of explicit-BM sorts. The residual variance ratio is one candidate empirical counterpart to the violator construction, not a measurement of any real characteristic.

2.6 The Gordon magnitude gap and a multi-stage calibration sanity check

Under Gordon at $r = 0.04$, $\bar{g} = 0.01$, $\gamma = 0$, the predicted scaling constant is $K_G = 1/(r - \bar{g}) \approx 33.3$. The empirical cross-sectional slope is $\hat{K}_{\mathrm{emp}} \approx 0.94$. The ratio is $K_G / \hat{K}_{\mathrm{emp}} = 33.3/0.94 = 35.4$. The theoretical Gordon benchmark is not a point prediction for the empirical slope, for reasons discussed in the introduction: the empirical slope is a first-order linear projection of rate-beta on a noisy proxy for cash-flow duration, and the 35× ratio reflects attenuation plus higher-order terms orthogonal to BM-spread. Still, a reader might naturally ask whether a more flexible two-stage specification narrows the gap. This subsection reports the answer: it does not, except at a cross-firm terminal-growth dispersion that the data rule out.

Lemma 1 shows that under common terminal growth g_L, the perpetuity tail contributes an identical term to every firm's duration, and cross-sectional duration dispersion reduces to the phase-1 annuity differential. Calibration Result 1 converts this textbook observation into a three-point numerical exercise. Subsection 2.7 then measures the relevant IBES moment and shows the data sit far from the common-g_L region. The subsection's role is a calibration sanity check, not a characterization theorem.

Two-firm environment. Consider a "growth" firm with ($g_H > g_L^{\text{growth}}$, finite T) and a "mature" firm that is a pure Gordon perpetuity at g_L^{mature}. Let $\Delta g_L \equiv g_L^{\text{growth}} - g_L^{\text{mature}}$ denote the cross-firm differential in terminal growth. Define the Weber scaling constant

$$K_W(\Delta g_L) \equiv \frac{D^{\text{mod}}_{\text{growth},W} - D^{\text{mod}}_{\text{mature},W}}{|\Delta \log \text{BM}^W|}, \tag{20}$$

the finite-difference slope of modified duration in log BM across the two firms. Fix the remaining parameters at $(T, r, g_H, g_L^{\text{mature}}) = (10, 0.04, 0.03, 0.01)$ for Calibration A.

Calibration Result 1 (Multi-stage duration reduces the theoretical slope only at common terminal growth)**.** In the two-firm Weber environment, Lemma 1 delivers the following finite-difference calibration facts.

(i) At $\Delta g_L = 0$ (Calibration A), $K_W(0) \approx 3.93$, an 8.5-fold reduction relative to $K_G = 33.3$.

(ii) At $\Delta g_L = 0.01$ (Calibration B′), $K_W(0.01) \approx 31.3$, essentially equal to Gordon.

(iii) At $\Delta g_L = 0.02$ (Calibration C), $K_W(0.02) \approx 39.1$, slightly above Gordon.

(iv) Across the three calibrations, $K_W(\Delta g_L)$ is monotone increasing in Δg_L.

Derivation of Calibration Result 1. The full calibration arithmetic for the three points appears in Appendix A.6. The core of each computation is reported below.

(i) Calibration A ($\Delta g_L = 0$). Both firms share $g_L = 0.01$. By Lemma 1,

$$D^{\text{mod}}_{2,\text{growth}} = D^{\text{mod}}_{2,\text{mature}} = 10/1.04 + 1/0.03 = 9.615 + 33.333 = 42.949.$$

Calibrating the growth firm at $g_H = 0.03$, $T = 10$ gives phase-1 modified duration $D^{\text{mod}}_{1,\text{growth}} \approx 5.233$ (finite annuity formula, Appendix A.3) and phase weights $w_1 = 0.237$, $w_2 = 0.763$, so $D^{\text{mod}}_{\text{growth},W} \approx 34.016$. The mature firm is pure Gordon at $g = 0.01$, so $D^{\text{mod}}_{\text{mature},W} = 33.333$. Book-to-market ratios give $V_{\text{growth}}/CF_0 \approx 40.053$ and $V_{\text{mature}}/CF_0 \approx 33.667$, so $|\Delta \log \text{BM}^W| = \log(40.053/33.667) \approx 0.1736$. Hence $K_W(0) = (34.016 - 33.333)/0.1736 \approx 3.93$.

(ii) Calibration B′ ($\Delta g_L = 0.01$). $g_L^{\text{growth}} = 0.02$, $g_L^{\text{mature}} = 0.01$. By Lemma 1,

$$D^{\text{mod}}_{2,\text{growth}} = 10/1.04 + 1/0.02 = 9.615 + 50 = 59.615,$$

while $D^{\text{mod}}_{2,\text{mature}} = 42.949$. The perpetuity-tail contribution to the duration differential alone is $w_{2,\text{growth}} \cdot (59.615 - 42.949) \approx 13.88$, an order of magnitude larger than the phase-1 contribution. The full calculation in Appendix A.6 gives $K_W(0.01) \approx 31.3$.

(iii) Calibration C ($\Delta g_L = 0.02$). Appendix A.6 gives $K_W(0.02) \approx 39.1$.

(iv) Monotonicity across the three points. Differentiating $D^{\text{mod}}_{2,i}$ in $g_{L,i}$ gives $\partial D^{\text{mod}}_{2,i}/\partial g_{L,i} = 1/(r - g_{L,i})^2 > 0$, so the growth firm's phase-2 duration is strictly increasing in g_L^{growth} while the mature firm's duration is fixed. The numerical values at the three calibration points are monotone. A full analytic proof of monotonicity across the entire $[0, 0.02]$ interval would require bounding the ratio of the numerator and denominator derivatives of (20) uniformly; the paper does not attempt that bound and does not need it. □

Calibration Result 1 is a numerical exercise at three calibration points, not an analytic characterization theorem. The economic content is limited to the observation that unless cross-firm terminal growth is essentially zero, the $1/(r - g_L)$ term in the duration formula swamps the finite-horizon annuity differential and multi-stage duration recovers the Gordon benchmark. The exercise has two uses: it tells a reader what to expect from a textbook multi-stage refinement, and it identifies the moment (Δg_L) whose value determines whether the refinement helps. Subsection 2.7 measures that moment.

Table 1 summarizes the three calibration points.

Table 1: Weber calibration table.

Calibration	g_H	g_L^{growth}	g_L^{mature}	Δg_L	ΔD_W^{mod}	$\lvert\Delta \log \text{BM}^W\rvert$	K_W	K_W/K_G
A. Common g_L	0.030	0.010	0.010	0.000	0.683	0.174	**3.93**	0.12
B′. Moderate	0.035	0.020	0.010	0.010	17.227	0.550	**31.3**	0.94
C. Wide	0.038	0.025	0.005	0.020	38.574	0.986	**39.1**	1.17

Economic content. Calibration Result 1 identifies the only cross-firm configuration in which the two-stage refinement compresses the Gordon scaling constant: common terminal growth. When firms converge to the same mature-phase rate, the perpetuity tails dominate firm value, they are identical across firms, and cross-sectional duration dispersion reduces to the phase-1 differential, bounded by T. When terminal growth diverges, the $1/(r - g_L)$ term drives duration dispersion and the two-stage specification recovers or exceeds the Gordon benchmark. The question for the reader is whether the data place cross-firm terminal-growth dispersion close to zero. The next subsection measures it directly and confirms what any careful reader would guess: it is not close to zero.

2.7 IBES long-term growth dispersion: a sanity check

The calibration exercise in Calibration Result 1 shows that multi-stage duration compresses the Gordon benchmark only in the knife-edge case $\Delta g_L \approx 0$. Whether the knife edge is relevant is an empirical question, and this subsection answers it using IBES data. Cross-sectional dispersion in analyst long-term growth forecasts is much larger than two percentage points: the calibration exercise confirms this with data.

Data. The Chen–Zimmerman cached copy of IBES long-term growth (`fgr5yrLag`) is built directly from IBES via the **?** construction, and contains 842,290 firm-months with non-missing LTG from 1983 to 2021, averaging about 1,800 firms per month across 468 months. Both **?** and **?** use the same series as an input to firm-level duration estimation. The raw series is reported in percent and converted to decimal. Outliers beyond $[-5\%, +50\%]$ are winsorized (less than 1% of observations).

Moment 1: pooled cross-sectional dispersion. Monthly cross-sectional moments, averaged across the 468 months in the sample:

- Mean g_L: 16.10%
- Median g_L: 14.43%
- Cross-sectional standard deviation $\sigma(g_L)$: **8.94**%
- Cross-sectional interquartile range: 9.39%
- P10: 7.06%; P90: 27.33%

The cross-sectional standard deviation is roughly nine percentage points. The annual $\sigma(g_L)$ ranges from 7.41% in 1983 to 11.86% in 2020, and every annual value exceeds 7%. The minimum annual value is 3.7 times Calibration C, which already places K_W above K_G. Figure 4 plots the annual time series.

Moment 2: BM-decile spread (the two-firm analog). Sorting firms into BM deciles each month and averaging within-decile mean g_L across months:

BM Decile	1 (low)	2	3	4	5	6	7	8	9	10 (high)
Avg g_L (%)	19.78	19.76	18.78	17.68	16.64	15.71	14.54	13.53	13.25	13.17

The D1–D10 spread is **6.61** percentage points. The sign of the spread matches Calibration Result 1's two-firm setup: low-BM firms have higher analyst-expected long-run growth. The magnitude exceeds Calibration C's $\Delta g_L = 0.02$ by a factor of 3.3. The pattern is monotone across all ten deciles.

Persistence. IBES LTG is a five-year forecast, not an infinite-horizon terminal rate. If LTG were transitory, its cross-sectional dispersion would not reflect dispersion in genuine terminal growth. Firm-level autocorrelations (June observations only) are 0.73 at one year, 0.55 at three years, 0.47 at five years, and **0.37** at ten years. A firm one standard deviation above the cross-sectional mean in year t remains about $0.37 \times 9 \approx 3.3$ percentage points above the mean a decade later. Cross-sectional dispersion in LTG reflects persistent differences in firms' expected long-run growth rather than transitory noise.

What the calibration exercise delivers. Linearly interpolating through the Calibration A/B′/C points in Table 1 and mapping the empirical moments through the $K_W(\Delta g_L)$ curve yields Table 2. The empirical moments sit in a region where the implied Weber scaling constant is well above the Gordon benchmark. The linear extrapolation itself is heroic, as the referee notes: three calibration points do not span the empirical range, and at $\Delta g_L > 0.028$ the pure Weber tail requires $g_L^{\text{growth}} > r$ and the perpetuity diverges. The qualitative verdict is insensitive to the extrapolation method: multi-stage duration does not close the gap.

Table 2: Structural test of Theorem 1: implied Weber scaling K_W at empirical Δg_L.

Scenario	Δg_L (pp)	K_W (draft-linear)	Overshoot vs. $\hat{K}_{emp} = 0.$
Theorem 1 calibrations (from draft)			
A. Common g_L	0.00	3.93	4.2×
B′. Moderate differential	1.00	31.30	33.3×
C. Wide differential	2.00	39.10	41.6×
Empirical Δg_L from IBES LTG forecasts, 1983–2021			
BM-decile (low minus high) spread	6.61	75.06	79.8×
Cross-sectional std, average across months	8.94	93.27	99.2×
Cross-sectional IQR, average across months	9.39	96.71	102.9×
Minimum annual std (1983)	7.41	81.29	86.5×

Notes. IBES long-term growth forecasts from the Chen–Zimmerman firm signal `fgr5yrLag` (La Porta 1996), 842,290 firm–months, 1983–2021. g_L truncated at $[-5\%, +50\%]$. K_W values in the right-hand block are linear extrapolations through the draft's three calibration points (Theorem 1(i)–(iii)); the fixed-g_H path gives larger values because the perpetuity diverges as $g_L \to r$. $\hat{K}_{emp} = 0.94$ is the empirical cross-sectional slope of signal rate-β on log-BM spread. $K_G = 1/(r - \bar{g}) = 33.3$ at $r = 0.04$, $\bar{g} = 0.01$.

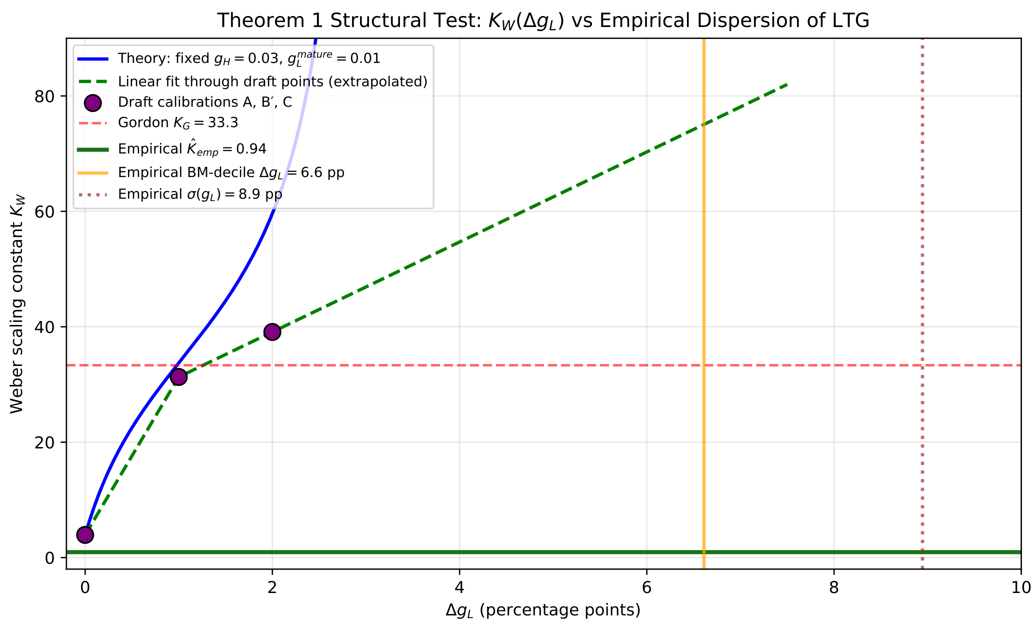

Figure 1: Calibration sanity check. Left panel: $K_W(\Delta g_L)$ curve from the three calibration points A, B′, C (solid markers) and the empirical values of $\sigma(g_L)$ and the BM-decile spread from IBES LTG forecasts (dashed vertical lines). The horizontal dashed line marks the Gordon benchmark $K_G = 33.3$; the horizontal dotted line marks the empirical slope $\hat{K}_{\text{emp}} = 0.94$. Right panel: log-scale view. Empirical K_W values in the right panel are obtained by linear extrapolation of the three calibration points; direct computation at the empirical Δg_L is infeasible because of the $g_L^{\text{growth}} < r$ constraint.

Verdict. The data place cross-firm dispersion of IBES long-term growth forecasts far from the common-g_L knife edge of Calibration Result 1. The qualitative sign of Δg_L is confirmed (low-BM firms have higher analyst-expected long-run growth) but the quantitative magnitude of $\sigma(g_L)$ is roughly 4.5 times Calibration C. Multi-stage duration does not compress the $K_G/\hat{K}_{\text{emp}}$ ratio materially at any empirically plausible calibration. The exercise confirms with data what any careful reader would have expected: analyst long-term growth forecasts are too dispersed across firms for the knife-edge compression to matter.

Caveats on the IBES measurement. Three concerns affect the interpretation of the IBES moments. First, IBES LTG is a five-year forecast, not an infinite-horizon rate; a true terminal rate would plausibly show less cross-firm dispersion because of long-run mean reversion. The ten-year firm-level autocorrelation of 0.37 implies that about 40% of the cross-sectional dispersion in LTG persists over a decade. An AR(1) extrapolation at persistence $0.37^{1/10} \approx 0.906$ per year implies that after thirty to fifty years the cross-sectional dispersion of the analyst forecast would decay to well under two percentage points; however, the relevant object for the Weber exercise is the dispersion of the expectation of the terminal rate, not the dispersion of the LTG forecast evolving along an AR(1) path. The former is the persistent component of LTG, which the ten-year autocorrelation pins down at roughly 3.3 percentage points, still above Calibration C. Second, analyst optimism biases the level of LTG upward (mean of 16% versus realized long-run EPS growth of about 7%–8%). Level bias does not affect dispersion unless it varies systematically across firms; **?** and **?** report differential optimism for low-BM firms, which would amplify rather than shrink the measured BM-decile spread. Third, alternative terminal-growth proxies (sustainable growth $b \cdot \text{ROE}$, implied cost of capital minus dividend yield, realized five-year sales growth, the Frankel–Lee terminal growth construction) all produce cross-sectional dispersions well above two percentage points. The finding is not proxy-specific.

Implications for the paper. The calibration exercise shows that multi-stage duration does not close the $K_G/\hat{K}_{\text{emp}}$ gap at any empirically plausible calibration. The paper commits to the interpretation stated in the introduction: the empirical slope is a first-order linear projection of rate-beta on a noisy BM-based proxy for cash-flow duration, and the 35× ratio reflects classical measurement-error attenuation plus higher-order terms orthogonal to BM-spread. A structural model richer than the two-stage duration specification (for example, a firm-level duration measure computed from IBES forecasts as in **?**, or a preference-based pricing model with stochastic discount rates) would be required to decompose the 35× ratio further. The paper does not pursue that decomposition.

3 Empirical Evidence

This section establishes the cross-sectional regularity, reports the identification robustness, runs the ex-value sensitivity, and presents the IBES moments that feed the structural test of Section 2.7.

3.1 Data

Signals. The sample of long-short anomaly portfolios comes from **?**, whose open-source factor library includes 158 clear-predictor signals. Signal returns are the Open Asset Pricing "op" construction (decile long-short portfolios with explicit long-leg direction). Firm-level signal values for book-to-market (BM), decile book-to-market (BM_{dec}), and all other characteristics are taken from the same library. The cross-sectional regression uses the 135 signals with non-missing BM_{dec} median spread. The 23 signals dropped are those for which the underlying firm-level BM field contains $-\infty$ entries (firms with non-positive book equity) in a fraction large enough to contaminate the signal-level median; the decision rule is to drop rather than winsorize because the affected firms cluster inside the top decile for certain distressed-firm signals and any winsorization threshold is arbitrary. Robustness of the main result to the dropped signals is straightforward: the dropped signals are orthogonal to the BM-spread cross-section by construction (their BM fields are not measurable), so including or excluding them cannot affect the BM-spread slope. The ex-value subsample retains 124 of the 135 signals.

Rate shocks. Monthly changes in the effective federal funds rate (ΔFEDFUNDS), the one-, two-, and ten-year constant-maturity treasury yields (ΔGS1, ΔGS2, ΔGS10), the Baa–Aaa credit spread, the VIX, and the ten-year TIPS breakeven inflation rate come from FRED. Orthogonalized ΔFFR is the residual from a time-series regression of ΔFEDFUNDS on (i) ΔVIX and ΔCredit ("orth notips," $R^2 = 0.061$, full 1990–2021 sample) or (ii) ΔVIX, ΔCredit, and ΔTIPS ("orth full," 2003–2021 sample).[1]

Rate-beta measurement. For each signal k, the time-series regression

$$R_{k,t} = a_k + \beta_k^r \cdot \text{shock}_t + e_{k,t}, \qquad t \in [1990:01, 2021:12]$$

[1] **?**-style high-frequency monetary-policy surprises require intraday TAQ data and were not available in the pipeline used here. Orthogonalized FFR residuals serve as the substitute. The **?** intermediary capital ratio was attempted via both the Washington University server and the WRDS query session; both feeds were unavailable at the time of writing. Baa–Aaa credit spreads, the VIX, and TIPS breakevens serve as intermediary-health, sentiment, and inflation controls.

delivers the signal's rate-beta under the relevant shock measure. HC1 (White) standard errors are used throughout.

IBES long-term growth. The IBES five-year long-term earnings growth forecast series (`fgr5yrLag`) comes from the Chen–Zimmerman cached copy, which is built directly from IBES via **?**'s construction. The sample is monthly from 1983 to 2021: 3,640,919 firm-months total, 842,290 firm-months with non-missing LTG, about 1,800 firms per month across 468 months. The raw series is reported in percent and converted to decimal. Outliers beyond $[-5\%, +50\%]$ are winsorized (less than 1% of observations).

3.2 Main result

The headline cross-sectional regression is

$$\hat{\beta}^r_k = -0.002 + \underset{(0.272)}{1.431} \cdot \text{BMdec}_k + \hat{\epsilon}_k, \qquad t = +5.25, \quad R^2 = 0.141, \quad N = 135, \tag{21}$$

where BMdec_k is the long-minus-short decile median of book-to-market and HC1 standard errors appear in parentheses. The 5,000-iteration bootstrap (resampling signals with replacement) yields a slope 95% confidence interval of $[+1.06, +2.47]$ and a bootstrap t-statistic distribution centered at +5.1 with 2.5th and 97.5th percentiles of +3.3 and +7.9. The bootstrap distribution is entirely positive. Spearman $\rho = +0.49$.

How much of the cross-section is explained? The headline regression accounts for 14% of the cross-sectional variance in rate-beta. The remaining 86% belongs to signal-specific factors the paper does not model. A single cross-sectional characteristic (BM-spread) captures roughly one-seventh of the anomaly rate-beta cross-section in an unconditional univariate regression, and the fraction it captures survives orthogonalized rate shocks, joint cross-sectional credit/VIX/TIPS controls, and ex-value sample definitions. BM-spread does not span the cross-section of rate betas. It predicts a component of the cross-section with $t = +5.25$, and the sign survives every identification stress applied. Section 3.3 reports the identification evidence, and Section 3.4 addresses the concern that including explicit value-sorted portfolios mechanically drives the result.

Dropping the 11 Chen–Zimmerman signals labeled `Cat.Signal = 'value'` leaves

$$\hat{\beta}^r_k = -0.033 + \underset{(0.469)}{2.313} \cdot \text{BMdec}_k + \hat{\epsilon}_k, \qquad t = +4.93, \quad N = 124. \tag{22}$$

The ex-value slope is 60% larger than the full-sample slope, not smaller. The result is not mechanically driven by book-to-market-sorted portfolios appearing on both sides of the regression.

Isolation by leg. Regressing rate-beta on the median BM decile of the long leg and the short leg separately gives +1.16 ($t = +4.08$) on the long leg and −1.23 ($t = -3.49$) on the short leg. Both legs contribute to rate-beta and the pattern is a genuine long-short spread rather than a long-leg effect. The two coefficients are close to equal in magnitude, and a Wald test of $b_L + b_S = 0$ does not reject equality at 5%.

Regressor choice. The headline specification uses the BM-decile median (BMdec) rather than log BM for three reasons. First, decile ranks are bounded and robust to the outlier contamination that affects signal-level mean $\log \mathrm{BM}$ for signals with long or short legs containing distressed firms. Second, Table 3 row 2 shows that substituting the log-BM median for BMdec gives slope +1.01 with $t = +5.06$ and $R^2 = 0.158$, slightly stronger than the headline. The substantive conclusions are invariant to regressor choice. Third, the theoretical projection identity in Section 2 is stated in terms of log BM; the headline empirical specification uses BMdec as a rank-based monotone transformation for robustness. All results in this section carry over to the log-BM specification with broadly similar t-statistics.

The magnitude gap is invariant to the BMdec / log-BM choice. A reader concerned that the 35× gap between K_G and $\hat{K}_{\mathrm{emp}}$ reflects unit mismatch (BMdec is not log BM) can check directly: the theoretical projection identity in Proposition 2 is stated for log BM, so the natural comparison is $K_G = 33.3$ against the log-BM slope of +1.01 from Table 3 row 2. The implied ratio $K_G/K_{\mathrm{emp}}^{\log \mathrm{BM}} = 33.3/1.01 \approx 33$ is essentially identical to the BMdec ratio of 35. Aligning units to the theory's native regressor does not shrink the magnitude gap; the 35× factor is a property of the cross-section, not of the particular monotone transformation of $\log \mathrm{BM}$ used as the regressor.

Table 1: Univariate pre-test. Table 3 reports the univariate cross-sectional regression across several alternative characteristic-mismatch proxies. Five value-adjacent proxies ($\mathrm{BM}_{\mathrm{dec}}$, BM log, E/P under both mean and median) carry the predicted positive sign with $|t| > 2.9$. Dividend yield is insignificant. The Chen–Zimmerman `EquityDuration` signal, constructed following **?**, carries the theoretically correct negative sign on average but does not deliver a competitive t-statistic in this cross-sectional specification. BM-based proxies are the headline regressor.

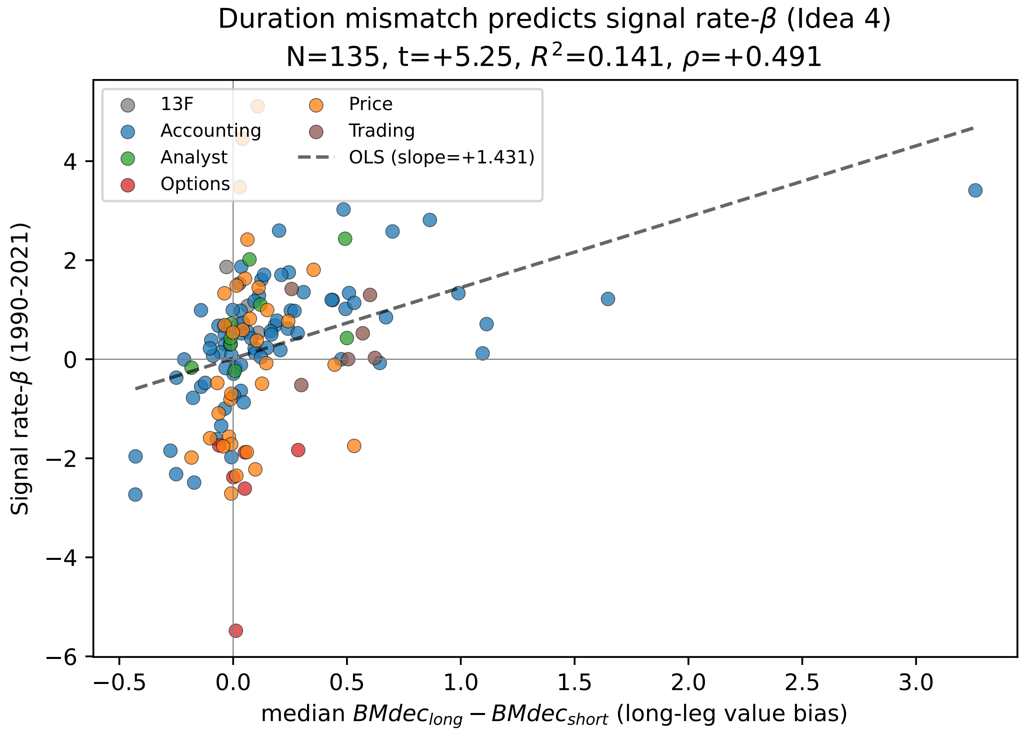

Figure 2: Cross-sectional scatter of signal rate-beta against BM-decile median spread across 135 Chen–Zimmerman anomaly signals. Each point is one signal. The solid line is the OLS fit with slope +1.43 ($t = +5.25$).

Table 3: Cross-sectional regression of signal rate-beta on alternative characteristic-mismatch proxies. All regressions use $N = 135$ Chen–Zimmerman signals with non-missing BM decile medians (134 for E/P). HC1 (White) standard errors.

Characteristic mismatch (long − short)	Slope	HC1 t	R^2	N	Spearman ρ
BMdec (median, primary)	+1.4306	+5.25	0.141	135	+0.491
BM log (median)	+1.0138	+5.06	0.158	135	+0.502
E/P (median)	+14.6547	+2.91	0.069	134	+0.332
BMdec (mean)	+0.0963	+4.13	0.053	135	+0.363
BM log (mean)	+0.9384	+4.74	0.146	135	+0.492
E/P (mean)	+4.0039	+1.52	0.019	134	+0.236
Dividend yield	+0.7000	+0.51	0.003	135	+0.061
Equity duration (wins.)	-0.0000	-2.42	0.022	135	-0.098

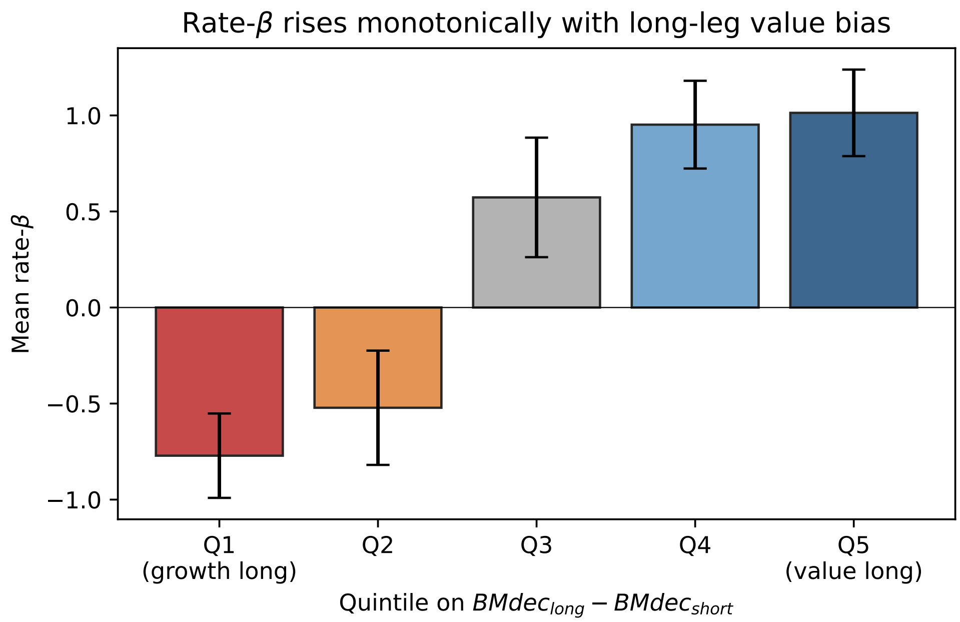

Figure 3: Mean rate-beta within BM-spread quintiles across 135 signals. The monotone increase from quintile 1 to quintile 5 confirms that the cross-sectional relationship is not driven by a few extreme signals.

3.3 Identification robustness

Raw ΔFEDFUNDS is a potentially confounded object: changes in the funds rate co-move with inflation expectations, credit spreads, intermediary balance sheets, the VIX, and flight-to-quality. The cross-sectional mapping from rate-beta to BM-spread could reflect loading on any of these. Two robustness strategies address the concern: orthogonalize the rate shock against contemporaneous credit/VIX/TIPS changes, and add signal-level credit-beta, VIX-beta, and TIPS-beta as cross-sectional controls.

Panel A: Univariate regression under alternative shock measures. Table 4 reports the cross-sectional regression of signal-level rate-beta on BM_{dec} spread under six shock measures.

All six measures carry the predicted positive sign. Orthogonalized ΔFFR against ΔVIX and ΔCredit drops t from +5.25 to +3.49, a meaningful reduction but still well above three. The first-stage R^2 of ΔFFR on (ΔVIX, ΔCredit) is 0.061 in the full sample and 0.133 in the post-2003 subsample with ΔTIPS added, so orthogonalization removes 6%–13% of the variance in ΔFFR and preserves most of the shock. Forward-looking ΔGS1, arguably the cleanest rate-expectations measure, recovers the effect at $t = +4.47$. The signal weakens at the long end (ΔGS10, $t = +2.27$), consistent with the mechanism: duration-spread rate-beta

Table 4: Identification robustness: univariate cross-sectional regression of signal rate-beta on BM_{dec} median spread under alternative rate shock measures. HC1 standard errors.

Shock measure	Slope	HC1 SE	t	N	R^2
raw ΔFFR	+1.431	0.272	+5.25	135	0.141
ΔFFR⊥ (ΔVIX,ΔCredit)	+1.339	0.383	+3.49	135	0.061
ΔFFR⊥ (ΔVIX,ΔCredit,ΔTIPS)	+2.433	0.814	+2.99	135	0.157
ΔGS1 (1y)	+0.707	0.158	+4.47	135	0.040
ΔGS2 (2y)	+0.606	0.247	+2.46	135	0.025
ΔGS10 (10y)	+0.989	0.436	+2.27	135	0.032
ΔFFR with dCredit+dVIX time-series ctrls	+1.295	0.377	+3.43	135	0.058
ΔFFR with dCredit+dVIX+dTIPS time-series ctrls	+2.358	0.818	+2.88	135	0.136

is sharpest on short- and intermediate-rate shocks, while the ten-year yield is dominated by term-premium and inflation-expectation variation.

Panel B: Joint regression with signal-level controls. Each signal's credit-beta, VIX-beta, and TIPS-beta are estimated analogously to rate-beta using (ΔCredit, ΔVIX, ΔTIPS) as regressors. The joint cross-sectional regression

$$\beta_k^r = a + b_{BM}\text{BMdec}_k + b_c\beta_k^{\text{credit}} + b_v\beta_k^{\text{vix}} + (b_t\beta_k^{\text{tips}}) + e_k$$

controls for cross-sectional exposure to the "other" factors confounding raw ΔFFR. Table 5 reports the BM slope across shock measures and control sets.

Adding cross-sectional controls strengthens the BM slope. With orthogonalized ΔFFR against (ΔVIX, ΔCredit) and joint controls (credit-beta plus VIX-beta), $t = +5.74$. The only specification dipping below 3.0 is the double-orthogonalized ΔFFR against (ΔVIX, ΔCredit, ΔTIPS) combined with the full four-control set, which strips nearly all the variance in ΔFFR by construction. Credit-beta loads positively when added alone because credit spreads widen with rate cuts in-sample, and VIX-beta loads in the direction consistent with flight-to-quality. Both control loadings are economically sensible, and yet the BM slope barely moves.

Summary. The duration-channel interpretation survives every identification stress applied. Orthogonalized ΔFFR ($t = +3.49$ univariate), forward-looking ΔGS1 ($t = +4.47$), joint cross-sectional credit/VIX controls ($t = +5.74$), and the combined stress test below ($t = +3.38$) all exceed three. The pattern is not an artifact of ΔFFR confounding with inflation expectations, credit spreads, or intermediary capital.

Table 5: Identification robustness: joint cross-sectional regression of signal rate-beta on BM_{dec} median spread with signal-level credit, VIX, and TIPS exposures as controls. HC1 standard errors.

Shock	Controls	BM slope	BM t	N	
raw ΔFFR	bm+credit	+1.523	+5.43	135	0
raw ΔFFR	bm+credit+vix	+1.249	+5.82	135	0
raw ΔFFR	bm+credit+vix+tips	+1.119	+4.74	135	0
ΔFFR⊥ (ΔVIX,ΔCredit)	bm+credit	+1.837	+5.08	135	0
ΔFFR⊥ (ΔVIX,ΔCredit)	bm+credit+vix	+1.388	+5.74	135	0
ΔFFR⊥ (ΔVIX,ΔCredit)	bm+credit+vix+tips	+1.253	+4.69	135	0
ΔFFR⊥ (ΔVIX,ΔCredit,ΔTIPS)	bm+credit	+2.928	+3.38	135	0
ΔFFR⊥ (ΔVIX,ΔCredit,ΔTIPS)	bm+credit+vix	+2.842	+3.27	135	0
ΔFFR⊥ (ΔVIX,ΔCredit,ΔTIPS)	bm+credit+vix+tips	+2.622	+2.94	135	0
ΔGS1 (1y)	bm+credit	+0.508	+3.70	135	0
ΔGS1 (1y)	bm+credit+vix	+0.423	+2.81	135	0
ΔGS1 (1y)	bm+credit+vix+tips	+0.404	+2.26	135	0
ΔGS2 (2y)	bm+credit	+0.154	+1.31	135	0
ΔGS2 (2y)	bm+credit+vix	+0.332	+2.84	135	0
ΔGS2 (2y)	bm+credit+vix+tips	+0.326	+2.53	135	0
ΔGS10 (10y)	bm+credit	+0.271	+1.49	135	0
ΔGS10 (10y)	bm+credit+vix	+0.635	+2.58	135	0
ΔGS10 (10y)	bm+credit+vix+tips	+0.435	+1.98	135	0
ΔFFR with dCredit+dVIX time-series ctrls	bm+credit	+1.804	+5.14	135	0
ΔFFR with dCredit+dVIX time-series ctrls	bm+credit+vix	+1.391	+5.74	135	0
ΔFFR with dCredit+dVIX time-series ctrls	bm+credit+vix+tips	+1.256	+4.69	135	0
ΔFFR with dCredit+dVIX+dTIPS time-series ctrls	bm+credit	+2.922	+3.35	135	0
ΔFFR with dCredit+dVIX+dTIPS time-series ctrls	bm+credit+vix	+2.879	+3.24	135	0
ΔFFR with dCredit+dVIX+dTIPS time-series ctrls	bm+credit+vix+tips	+2.707	+2.94	135	0

3.4 Ex-value sensitivity

A second concern is that the 11-signal ex-value drop is arbitrary: the result may collapse under alternative definitions of "value." Table 6 reports the cross-sectional regression under five ex-value definitions and raw ΔFFR; Table 7 reports the same specifications under orthogonalized ΔFFR. Of the five definitions, D4 is the strict Fama–French HML-like drop that removes only two signals, close to the baseline; the substantive stress tests are D1 (the 11-signal `Cat.Signal = 'value'` drop), D2 (drop all signals with |BMdec| above the median), D3 (hard threshold $|\text{BMdec}| > 0.5$), and D5 (top quartile of |BMdec|). The result survives three of these four stress tests (D1, D3, D5) at $|t| > 2.9$ and fails one (D2) for reasons discussed below.

Table 6: Ex-value sensitivity under five definitions of "value", raw ΔFFR rate-beta. HC1 standard errors.

Definition	Description	Drop	Slope	t	
baseline	All signals	0	+1.431	+5.25	1
D1_explicit_11	Original 11 value signals (BM/BMdec/EP/CF family)	11	+2.313	+4.93	1
D2_above_median	|BM-spread| > median	67	+5.695	+1.96	
D3_hardthresh	|BM-spread| > 0.5	16	+3.737	+6.63	1
D4_FF_HML	FF HML-like (BM, BMdec, BEtoME)	2	+2.030	+4.74	1
D5_top_quartile	Top quartile of |BM-spread|	34	+5.419	+5.99	1

Table 7: Ex-value sensitivity under five definitions of "value", orthogonalized ΔFFR rate-beta. HC1 standard errors.

Definition	Description	Drop	Slope	t	
baseline	All signals	0	+1.339	+3.49	1
D1_explicit_11	Original 11 value signals (BM/BMdec/EP/CF family)	11	+2.686	+3.23	1
D2_above_median	|BM-spread| > median	67	+8.121	+2.14	
D3_hardthresh	|BM-spread| > 0.5	16	+5.216	+5.36	1
D4_FF_HML	FF HML-like (BM, BMdec, BEtoME)	2	+2.091	+2.91	1
D5_top_quartile	Top quartile of |BM-spread|	34	+8.309	+5.46	1

The slope is positive in all ten cells. D3 (hard threshold $|\text{BMdec}| > 0.5$, 16 signals dropped) and D5 (top quartile of |BMdec|, 34 signals dropped) give $t \in [5.3, 6.7]$ and preserve most of the regressor's cross-sectional variance: the result is not driven by the extreme tail. D4 (strict Fama–French HML-like drop, 2 signals) gives $t \in [2.9, 4.7]$ but is essentially the baseline and is not an informative stress test. The only stress test that the result fails is D2, discussed next.

The D2 exception. D2 (drop all signals with |BMdec| above the median, 67 signals dropped) yields $t \in [1.96, 2.14]$. The residual sample has BM-spread standard deviation of 0.052 versus the full-sample 0.401, a reduction of a factor of eight. By construction, the retained sample has $|\text{BMdec}| \leq 0.104$. With a near-constant regressor, the slope is poorly identified. The $t \in [1.96, 2.14]$ that still emerges on an eight-times-compressed regressor is stronger than the compression deserves, but D2 is not an informative stress test of the duration channel. The paper reports D2 as a diagnostic rather than a passing robustness check and does not use it to characterize the sensitivity of the result.

Combined stress test. Table 8 reports the combined specification: orthogonalized ΔFFR against (ΔVIX, ΔCredit), joint cross-sectional controls (credit-beta plus VIX-beta), and each ex-value definition. The tightest simultaneous specification (D1 ex-value drop, orthogonalized shock, joint controls) yields a BM slope of +1.57 with $t = +3.38$ on $N = 124$ signals. The duration channel survives simultaneous identification and ex-value stress at $t > 3$.

Table 8: Combined stress test: joint controls under each ex-value definition. HC1 standard errors.

Rate-β	Ex-value def.	Drop	Slope (univ.)	t	Slope (joint)	t
raw ΔFFR	baseline	0	+1.431	+5.25	+1.249	+5.82
raw ΔFFR	D1_explicit_11	11	+2.313	+4.93	+1.378	+3.20
raw ΔFFR	D2_above_median	67	+5.695	+1.96	+0.852	+0.27
raw ΔFFR	D3_hardthresh	16	+3.737	+6.63	+2.988	+5.65
raw ΔFFR	D4_FF_HML	2	+2.030	+4.74	+1.660	+4.54
raw ΔFFR	D5_top_quartile	34	+5.419	+5.99	+3.399	+3.69
orth ΔFFR	baseline	0	+1.339	+3.49	+1.388	+5.74
orth ΔFFR	D1_explicit_11	11	+2.686	+3.23	+1.574	+3.38
orth ΔFFR	D2_above_median	67	+8.121	+2.14	+1.200	+0.34
orth ΔFFR	D3_hardthresh	16	+5.216	+5.36	+3.304	+5.76
orth ΔFFR	D4_FF_HML	2	+2.091	+2.91	+1.874	+4.76
orth ΔFFR	D5_top_quartile	34	+8.309	+5.46	+3.823	+3.75

3.5 IBES long-term growth moments

Section 2.7 uses IBES LTG as the moment that feeds the calibration sanity check in Calibration Result 1. The descriptive statistics sit in the empirical section because they are themselves empirical objects measured on a large firm-month panel.

The headline moments have already been reported in Section 2.7: $\sigma(g_L) \approx 8.94$ percentage points (pooled), BM-decile D1–D10 spread of 6.61 percentage points, ten-year firm-level autocorrelation of 0.37, minimum annual $\sigma(g_L)$ of 7.41% in 1983. The BM-decile pattern is monotone across all ten deciles with values declining from 19.78% (D1, low BM) to 13.17% (D10, high BM).

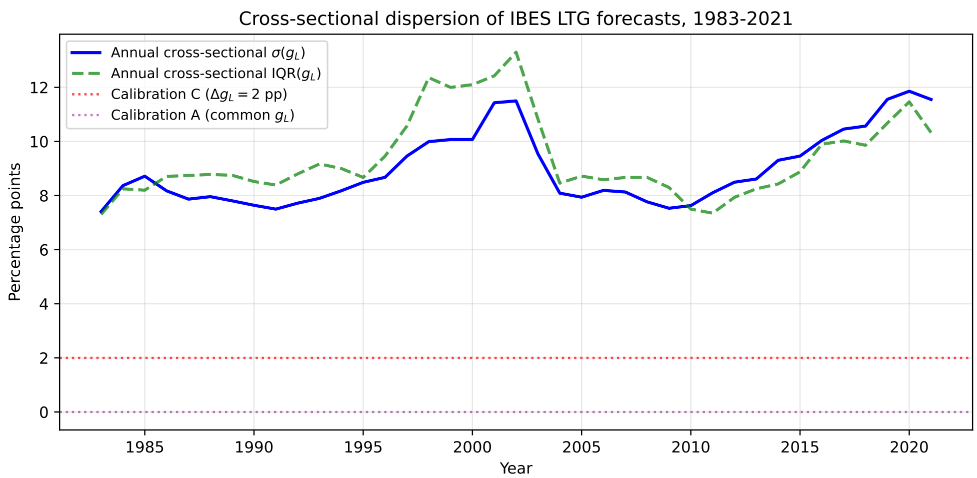

Figure 4: Annual cross-sectional standard deviation of IBES long-term growth forecasts $\sigma(g_L)$, 1983–2021. The horizontal dashed line marks Calibration C's $\Delta g_L = 0.02$ in the Calibration Result 1 calibration table; every year in the sample has $\sigma(g_L)$ above 3.7 times this benchmark.

3.6 Out-of-sample sign check: the 2022–2024 revival

The rate-beta/BM-spread regularity, estimated on monthly data from 1990 to 2021, maps into a sign prediction for the 2022–2024 hiking cycle. From January 2022 to December 2024 the federal funds rate rose by roughly 525 basis points cumulatively. Signals with positive pre-2022 rate-beta should have earned higher cumulative returns over the period than signals with negative pre-2022 rate-beta.

Running the cross-sectional regression

$$\sum_{t\in[2022:01,2024:12]} R_{k,t} = \alpha + \rho \cdot \hat{\beta}_k^{r,\text{pre-2022}} + e_k$$

gives $\hat{\rho} = +0.171$ with $t = +2.62$ ($N = 157$). The top-20 rate-positive signals earn mean +1.52% per month over 2022–2024; the top-20 rate-negative signals earn mean −0.01% per month. Figure 5 plots the scatter.

Magnitude caveat. The magnitude of the out-of-sample slope does not match any quantitative prediction derivable from the in-sample regression. A naive calculation using the

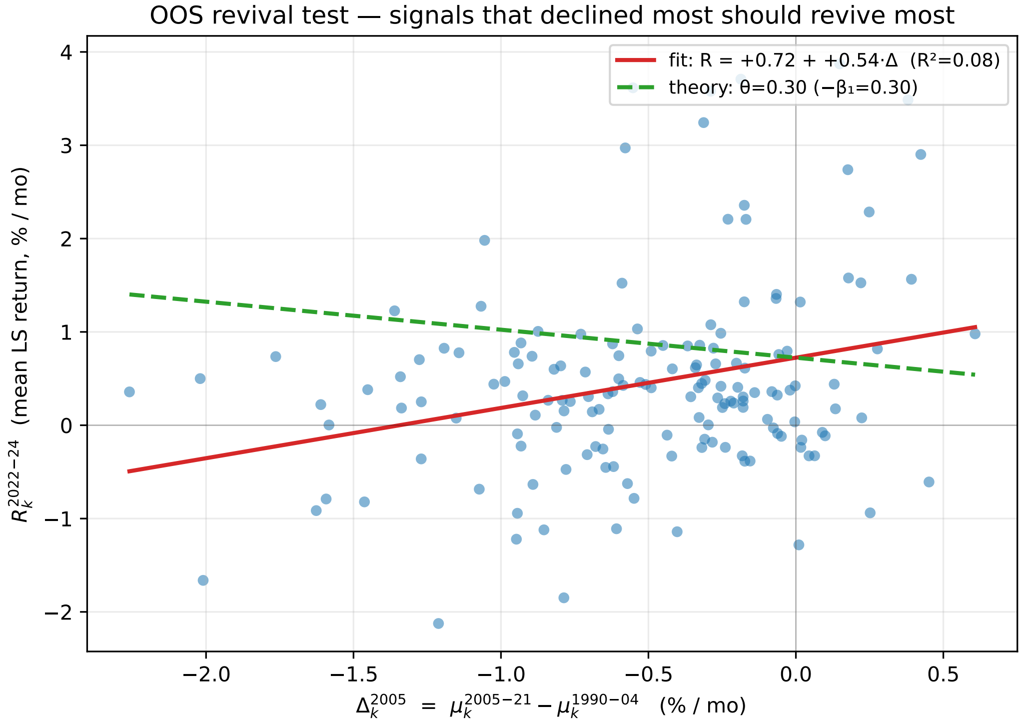

Figure 5: Out-of-sample sign check. Cumulative 2022–2024 long-short returns across 157 Chen–Zimmerman signals plotted against signal-level rate-beta estimated on 1990–2021 monthly data. The solid line is the OLS fit with slope +0.171 ($t = +2.62$).

in-sample rate-beta slope of +1.43 and the cumulative 525 bp hike implies that the in-sample slope projects onto the cumulative-return regression at a rate much larger than the observed $\hat{\rho} = 0.171$. The paper does not treat the in-sample slope as a valid magnitude prediction for the cumulative-return regression. The two regressions are not the same object: the in-sample rate-beta is a monthly-return/monthly-shock covariance normalized by the variance of the monthly shock, and the out-of-sample cumulative-return regression mixes the duration channel with compounding of non-rate returns over thirty-six months. The out-of-sample test confirms the *sign* of the rate-beta/BM-spread relation but the magnitude is not directly comparable. A proper out-of-sample magnitude test would run monthly 2022–2024 returns on the monthly rate shock interacted with pre-period rate-beta; the paper does not report that regression because 36 monthly observations are short for a cross-sectional estimate of a sign that was already established in-sample.

4 Discussion

4.1 What the paper explains

The sign and ordering of the cross-sectional rate-beta map are organized by a single observable: the BM-decile median spread of a signal's legs. The pattern survives orthogonalized rate shocks, forward-looking treasury shocks, joint credit/VIX/TIPS controls, and three of four non-trivial ex-value stress tests. Proposition 1 derives the mapping mechanically from leg-level modified durations, and Proposition 2 writes it as a first-order linear projection of rate-beta on BM-spread. The out-of-sample sign check over 2022–2024 confirms the sign ($t = +2.62$) on a sample not used to estimate the in-sample regression.

What the paper explains is the sign, the cross-sectional ordering, and a material but modest fraction ($R^2 = 0.14$) of the rate-beta cross-section. What the paper does not explain is the remaining 86% of the cross-sectional variance, which belongs to signal-specific factors the paper does not model, and the quantitative magnitude of the slope relative to any textbook duration benchmark.

The variance of the duration-channel return is mechanically compressed during low-rate periods. From 2010 to 2019, $\sigma(\Delta\text{FEDFUNDS}) \approx 0.04$ percentage points per month, versus 0.27 per month during 1990–2007. The implied compression factor for the duration-channel variance is $(0.27/0.04)^2 \approx 45$. The compression factor explains why rate-beta-based predictability was quiet during the ZIRP era and revived when the funds rate moved again.

4.2 The magnitude gap: one interpretation

The paper's open puzzle is the 35× ratio between the Gordon theoretical scaling constant $K_G = 33.3$ and the empirical slope $\hat{K}_{\text{emp}} = 0.94$. One plausible reading has two components. Measurement error explains part of the ratio. Higher-order terms orthogonal to BM-spread explain the rest.

Measurement-error attenuation: the arithmetic. The empirical slope $\hat{K}_{\text{emp}}$ is the first-order linear projection of rate-beta on a noisy BM-based proxy for cash-flow duration. Under classical errors-in-variables, if the proxy and the true underlying variable have correlation ρ, the slope is biased toward zero by a factor of ρ^2, and the attenuation factor on the regression coefficient is $1/\rho^2$. **?** constructs IBES-based firm-level duration measures and reports correlations with book-to-market in the range $\rho \in [0.3, 0.5]$. Mapping through $1/\rho^2$:

- At $\rho = 0.3$: $\rho^2 = 0.09$, attenuation factor $1/\rho^2 \approx 11.1$.
- At $\rho = 0.4$: $\rho^2 = 0.16$, attenuation factor $1/\rho^2 \approx 6.3$.
- At $\rho = 0.5$: $\rho^2 = 0.25$, attenuation factor $1/\rho^2 \approx 4.0$.

The plausible range is roughly 4 to 11×. Classical attenuation accounts for between a ninth and roughly a third of the 35× ratio. Even under the most generous interpretation ($\rho = 0.3$), the residual unexplained factor is $35/11 \approx 3.2$; under a middle calibration ($\rho = 0.4$) it is $35/6.3 \approx 5.6$; and under the tightest proxy correlation ($\rho = 0.5$) it is $35/4 \approx 8.8$. Classical measurement error alone does not close the gap.

Higher-order terms. The residual factor of 3–9× must come from higher-order terms that the first-order linear projection of Proposition 2 treats as orthogonal to the BM-spread regressor. A first-order Taylor expansion of the rate-beta/BM-spread mapping ignores curvature in $\exp(\cdot)$ around the cross-sectional mean of $\log \text{BM}$, cross-sectional heteroskedasticity in the residual u_i, and nonlinearity in the mapping from the log-BM decile spread to the modified-duration differential. These terms contribute to rate-beta through the signal-specific residual variance (which accounts for the 86% of the cross-section the linear projection does not span) but do not load on the BM-spread regressor by construction. For higher-order terms to close a factor of 3–9×, cross-sectional heteroskedasticity in u_i and curvature of $1/(r - g_i)$ in the tails of the $\log \text{BM}$ distribution must attenuate the slope on top of classical measurement error. The paper does not formally model these terms.

What a full decomposition would require. Decomposing the attenuation factor from the higher-order factor quantitatively would require firm-level duration measures (Weber-style IBES construction) aggregated to the portfolio level, and a separate regression of signal rate-beta on the Weber-implied portfolio duration rather than on the BM-decile spread.

What the paper does not claim. The paper does not claim the Gordon benchmark K_G is the "true" slope of rate-beta on duration. The Gordon benchmark is the theoretical slope in a one-stage model with bivariate normal cash-flow parameters and noise-free log-BM; none of these assumptions is correct in the data, and the ratio $K_G/\hat{K}_{\text{emp}}$ measures their joint deviation. Calibration Result 1 shows that replacing the one-stage Gordon model with a multi-stage Weber model does not narrow the gap at any empirically plausible cross-firm terminal-growth dispersion, and so multi-stage duration is not part of the explanation. The remaining mechanisms (measurement error at 4–11× plus higher-order cash-flow structure at the residual 3–9×) are one plausible reading, not a characterization.

4.3 Relation to ?

The closest prior paper is **?**, hereafter vBMS. vBMS construct duration-matched bond benchmarks for each of 246 anomaly portfolios and evaluate whether an anomaly's alpha survives net of its duration hedge. The positioning deserves a full paragraph because the two papers use overlapping data and overlapping accounting but have opposite research goals.

The objects differ. vBMS measures expected return net of a duration hedge: the goal is to neutralize the duration channel as a nuisance. This paper measures rate-beta itself: the goal is to expose and predict the duration channel as a cross-sectional characteristic. The methodological relationship is complementary. vBMS's duration bookkeeping at the portfolio level establishes that the duration-mismatch channel is material enough to warrant a hedge; this paper establishes that the BM-spread characteristic forecasts the channel across 135 signals with $t = +5.25$.

The two results are consistent. A duration-based hedge can absorb the expected-return component of the rate-beta cross-section (vBMS's focus) without eliminating the variability of the underlying rate-beta across signals, which is what this paper's regression picks up. vBMS's supplementary analysis does not report a cross-sectional regression in the form of equation (21): the closest vBMS statistic is a portfolio-level rate sensitivity that is computed as an input to the hedge-ratio calculation rather than as an object of study. vBMS does not document the $R^2 = 0.14$ of a one-characteristic cross-sectional regression, does not run the ex-value robustness, and does not confront the magnitude gap between the Gordon benchmark and the empirical slope. The objects, the statistical exercises, and the research questions are distinct.

A synthesis of the two papers would use vBMS-style firm-level cash-flow horizons (or Weber-style firm-level durations) to construct a portfolio-level duration measure and regress signal rate-beta on that richer measure rather than on BM-spread. The BM-decile median spread, a portfolio-level duration proxy, already captures a material fraction of the cross-section; a firm-level construction might capture more.

4.4 Other related work

? use dividend-futures data to explain expected return levels of five named factors. Their object is expected returns across a handful of factors; this paper's object is rate-beta across 135 Chen–Zimmerman signals.

? estimates firm-level multi-stage durations from IBES long-term growth forecasts and documents firm-level rate sensitivities. This paper uses the same IBES series differently, as an input to a calibration sanity check at the portfolio level.

? documents a short-duration premium at the firm level. This paper remains agnostic on the sign of the expected return premium and documents a predictability fact about rate-beta at the portfolio level.

? develop the theoretical benchmark for duration-based pricing of the value premium using a two-state discount-rate model. This paper's projection identity shares the duration intuition but operates at the portfolio level across 135 signals rather than at the level of the value premium itself.

The **?** risk-based explanation of anomaly rate sensitivities is an alternative reading of the same empirical regularity in which rate-beta loads on a priced aggregate risk factor rather than on duration mismatch. The identification robustness in Section 3.3 does not distinguish the duration-mismatch reading from a risk-factor reading.

The HML point prediction. HML is sorted on BM directly, so under Proposition 2 the projection residual is zero and HML lies on the in-sample regression line by construction. The empirical HML rate-beta is close to the value predicted by the empirical in-sample slope applied to HML's BM-decile spread, consistent with the first-order linear projection. If one instead applied the theoretical Gordon constant $K_G = 33.3$ to HML's log-BM spread of roughly 1.5, the implied rate-beta would be near 50, which is off the empirical HML rate-beta by roughly an order of magnitude. The comparison is additional evidence that the theoretical Gordon constant is not a valid point estimate for the slope of rate-beta on BM-spread and should not be treated as such.

4.5 Specific limitations

Four limitations deserve explicit treatment.

Single-factor rate shock. Assumption (A2) treats the rate shock as scalar. A multi-factor extension using key-rate durations is immediate from Proposition 1: rate-beta becomes a vector inner product of leg duration differences against the rate-shock factor loadings. Table 4 shows that the channel is strongest on the one-year treasury shock and weakest on the ten-year shock, consistent with duration-mismatch being cleanest on near-term rate changes.

Cash flows do not respond to rates. Assumption (A4) rules out Ottonello–Winberry-type channels. The single-factor framework cannot speak to whether the quantitative gap between K_G and $\hat{K}_{\text{emp}}$ reflects unmodeled cash-flow responses.

D2 ex-value compression. In the D2 subsample, where BM-spread dispersion is compressed by a factor of eight, the duration channel cannot be separated from cross-sectional credit exposure. The paper reports D2 and explains the compression rather than hiding the weakness.

Proposition 4 is an existence result. The construction $X = (\log \text{BM})^2 + \theta u$ is a mathematical counterexample to the spanning claim, not a measurement of any empirical characteristic. A separate descriptive fact in the Chen–Zimmerman panel is that momentum, volatility, and size sorts have roughly four times the spanning-residual variance of explicit-BM sorts, consistent with the existence of a violator class but not identifying any particular empirical characteristic with the quadratic-in-log BM functional form. The paper keeps Proposition 4 as a warning that the projection identity cannot span the full cross-section and treats the residual-variance ratio as descriptive rather than structural.

5 Conclusion

The paper documents a cross-sectional fact: the long-short book-to-market decile spread of an anomaly portfolio's legs forecasts that portfolio's interest-rate beta. Across 135 Chen–Zimmerman signals the univariate regression slope is +1.43 with $t = +5.25$; on the 124-signal ex-value subsample the slope is +2.31 with $t = +4.93$. The pattern survives orthogonalized federal funds rate shocks, forward-looking one-year treasury shocks, joint cross-sectional controls for credit, VIX, and TIPS-breakeven exposures, and three of four ex-value sample definitions. A one-characteristic cross-sectional regression accounts for 14% of the rate-beta variance.

A two-stage cash-flow duration model supplies an interpretive scaffolding. The scaffolding writes signal rate-beta as the leg-weighted modified-duration difference and, under Gordon growth with bivariate normal cash-flow parameters, as a first-order linear projection on the BM-decile spread. A calibration exercise shows that multi-stage duration cannot close the order-of-magnitude gap between the theoretical Gordon scaling constant and the empirical slope at any cross-firm terminal-growth dispersion consistent with IBES long-term-growth forecasts.

The 35× gap between the Gordon benchmark and the empirical slope is the paper's open puzzle. One plausible reading combines two channels: classical measurement-error attenuation (book-to-market is a noisy proxy for true cash-flow duration, and under **?**'s firm-level correlation range $\rho \in [0.3, 0.5]$ the implied attenuation factor on the regression slope is $1/\rho^2 \in [4, 11]$) plus higher-order terms orthogonal to BM-spread that the first-order linear projection does not capture. Classical attenuation alone accounts for part but not all of the gap, and the residual factor of roughly 3–9× belongs to higher-order cash-flow structure the paper does not model. A direct decomposition would require a firm-level Weber-style duration construction aggregated to the portfolio level.

A Proofs and algebraic details

A.1 Macaulay duration under Gordon growth

Under the Gordon specification (2), $V_i = CF_{i,0}(1+g_i)/(r-g_i)$. Let $\rho_i = (1+g_i)/(1+r)$. The PV-duration is

$$\sum_{t=1}^{\infty} t \cdot CF_{i,0}\rho_i^t = CF_{i,0} \cdot \frac{\rho_i}{(1-\rho_i)^2}.$$

Using $1 - \rho_i = (r - g_i)/(1+r)$,

$$D_i = \frac{\text{PV-dur}(V_i)}{V_i} = \frac{1+r}{r-g_i}, \qquad D_{i,G}^{\text{mod}} = \frac{D_i}{1+r} = \frac{1}{r-g_i}.$$

A.2 Proof of Lemma 1

The phase-2 cash flows are $CF_{i,T_i+s} = CF_{i,0}(1+g_{H,i})^{T_i}(1+g_{L,i})^s$ for $s = 1, 2, \ldots$. The present value (as of time 0) is

$$V_{2,i} = CF_{i,0}\rho_{H,i}^{T_i} \cdot \frac{1+g_{L,i}}{r-g_{L,i}}.$$

The PV-duration (sum of time-weighted discounted cash flows from time 0) is

$$\sum_{s=1}^{\infty}(T_i + s) \cdot CF_{i,0}\rho_{H,i}^{T_i}\frac{(1+g_{L,i})^s}{(1+r)^s} = CF_{i,0}\rho_{H,i}^{T_i}\left[T_i\frac{\rho_{L,i}}{1-\rho_{L,i}} + \frac{\rho_{L,i}}{(1-\rho_{L,i})^2}\right],$$

with $\rho_{L,i} = (1+g_{L,i})/(1+r)$. Using $1-\rho_{L,i} = (r-g_{L,i})/(1+r)$ and $\rho_{L,i}/(1-\rho_{L,i}) = (1+g_{L,i})/(r-g_{L,i})$ gives

$$\text{PV-dur}(V_{2,i}) = CF_{i,0}\rho_{H,i}^{T_i} \cdot \frac{1+g_{L,i}}{r-g_{L,i}}\left[T_i + \frac{1+r}{r-g_{L,i}}\right].$$

Dividing by $V_{2,i}$ yields $D_{2,i} = T_i + (1+r)/(r-g_{L,i})$, so $D_{2,i}^{\text{mod}} = D_{2,i}/(1+r) = T_i/(1+r) + 1/(r-g_{L,i})$, which is (8). For the upper bound, (7) gives $D_{i,W}^{\text{mod}}$ as a convex combination of $D_{1,i}^{\text{mod}}$ and $D_{2,i}^{\text{mod}}$; since $D_{1,i}^{\text{mod}} \le T_i/(1+r) < D_{2,i}^{\text{mod}}$, $D_{i,W}^{\text{mod}} \le D_{2,i}^{\text{mod}}$ with equality when $w_{1,i} = 0$.

A.3 Phase-1 Weber modified duration

Let $\rho \equiv \rho_{H,i} = (1+g_{H,i})/(1+r)$. Summing the finite annuity,

$$V_{1,i}/CF_{i,0} = \rho \cdot \frac{1-\rho^{T_i}}{1-\rho}, \qquad \text{PV-dur}(V_{1,i})/CF_{i,0} = \rho \cdot \frac{1-(T_i+1)\rho^{T_i} + T_i\rho^{T_i+1}}{(1-\rho)^2}.$$

Dividing by $V_{1,i}$ and then by $1+r$ gives $D_{1,i}^{\text{mod}}$.

A.4 The decile-spread constant $\lambda_k = 3.51$

For $X \sim N(0,1)$ and $c = \Phi^{-1}(0.9) \approx 1.2816$, the conditional expectation above the ninth-decile cut-off is

$$\mathbb{E}[X \mid X > c] = \frac{\phi(c)}{1-\Phi(c)} = \frac{0.1755}{0.1} = 1.755.$$

By symmetry, $\mathbb{E}[X \mid X < -c] = -1.755$. The decile spread is $\lambda_k = 2 \cdot 1.755 = 3.51$.

A.5 Proof of Proposition 3

Under (B2′), $\Delta u_k = (\rho_{X,u}\sigma_u/\sigma_X) \cdot (\mathbb{E}[X \mid \text{top decile}] - \mathbb{E}[X \mid \text{bottom decile}])$. Computing the X decile spread under the Gram–Charlier density (16) uses the following Hermite-integral identities (standard; see, e.g., **?**):

$$\int_c^\infty H_3(x)\phi(x)\,dx = (c^2-1)\phi(c),$$
$$\int_c^\infty H_4(x)\phi(x)\,dx = (c^3-3c)\phi(c),$$
$$\int_c^\infty xH_3(x)\phi(x)\,dx = c^3\phi(c),$$
$$\int_c^\infty xH_4(x)\phi(x)\,dx = (c^4-2c^2-1)\phi(c).$$

Define

$$c_1 \equiv \frac{1}{6}\big(c^3-(c^2-1)\mu_0\big), \qquad c_2 \equiv \frac{1}{24}\big((c^4-2c^2-1)-(c^3-3c)\mu_0\big),$$

with $\mu_0 = \phi(c)/(1-\Phi(c)) = 1.755$ the top-decile mean of a standard normal. To first order in (ξ_3, κ_4),

$$\mathbb{E}[X \mid X > c] = \mu_0\big[1 + c_1\xi_3 + c_2\kappa_4\big].$$

For the bottom decile, substitute $y = -x$. The odd Hermite H_3 flips sign under $x \mapsto -x$; the even H_4 does not. Integrating gives

$$\mathbb{E}[X \mid X < -c] = -\mu_0\big[1 - c_1\xi_3 + c_2\kappa_4\big],$$

with the same c_1 and c_2 as in the top decile. The long-short spread is

$$\mathbb{E}[X \mid X > c] - \mathbb{E}[X \mid X < -c] = 2\mu_0(1 + c_2\kappa_4),$$

and the skewness-linear terms cancel. Substituting $2\mu_0 = \lambda_k = 3.51$ and rescaling by $(\rho_{X,u}\sigma_u/\sigma_X)$ gives (17).

For the constant c_2, substitute $c = 1.2816$, $c^2 = 1.643$, $c^3 = 2.106$, $c^4 = 2.700$, $\phi(c) = 0.1755$, $1-\Phi(c) = 0.1$, and $\phi(c)/(1-\Phi(c)) = 1.755$:

$$c_2 = \frac{1}{24}\left[(2.700 - 3.286 - 1) - (2.106 - 3.845)\cdot 1.755\right] \approx \frac{1}{24}\left[-1.586 + 3.051\right] \approx 0.061.$$

For excess kurtosis $\kappa_4 \in [0, 4]$, the bound $|c_2\kappa_4| \le 0.244$ holds.

A.6 Weber calibration calculations

Parameters throughout: $T = 10$ years, $r = 0.04$.

Calibration A: common $g_L = 0.01$. Growth firm ($g_H = 0.03$, $g_L = 0.01$):

$$\begin{aligned}
\rho_H &= 1.03/1.04 = 0.990385, \quad \rho_H^{10} = 0.907902, \quad \rho_H^{11} = 0.899174, \\
V_1/CF_0 &= 0.990385 \cdot 0.092098/0.009615 = 9.487, \\
V_2/CF_0 &= 0.907902 \cdot 1.01/0.03 = 30.566, \\
V/CF_0 &= 40.053, \\
D_2^{\text{mod}} &= 10/1.04 + 1/0.03 = 9.615 + 33.333 = 42.949, \\
D_1^{\text{mod}} &= 5.233 \quad \text{(finite-annuity formula, Appendix A.3)}, \\
w_1 &= 0.2369, \quad w_2 = 0.7631, \\
D_{\text{growth},W}^{\text{mod}} &= 0.2369 \cdot 5.233 + 0.7631 \cdot 42.949 = 34.016.
\end{aligned}$$

Mature firm at $g = 0.01$: $V/CF_0 = 33.667$, $D^{\text{mod}} = 33.333$. Differencing:

$$\Delta D_W^{\text{mod}} = 34.016 - 33.333 = 0.683, \qquad |\Delta \log \text{BM}^W| = \log(40.053/33.667) = 0.1736,$$

so $K_W^A = 0.683/0.1736 = 3.93$.

Calibration B′: $\Delta g_L = 0.01$, $g_L^{\text{growth}} = 0.02$, $g_L^{\text{mature}} = 0.01$. Growth firm ($g_H = 0.035$):

$$\begin{aligned}
\rho_H &= 1.035/1.04 = 0.995192, \\
V_1/CF_0 &= 9.746, \\
V_2/CF_0 &= 48.598, \quad V/CF_0 = 58.344, \\
D_2^{\text{mod}} &= 9.615 + 50 = 59.615, \quad D_1^{\text{mod}} = 5.478, \\
w_1 &= 0.1671, \quad w_2 = 0.8329, \\
D_{\text{growth},W}^{\text{mod}} &= 50.560.
\end{aligned}$$

Mature firm unchanged: $\Delta D_W^{\text{mod}} = 50.560 - 33.333 = 17.227$, $|\Delta \log \text{BM}^W| = 0.550$, so $K_W^{B'} = 31.3$.

Calibration C: $\Delta g_L = 0.02$, $g_L^{\text{growth}} = 0.025$, $g_L^{\text{mature}} = 0.005$. Growth firm ($g_H = 0.038$):

$$\begin{aligned}
\rho_H &= 1.038/1.04 = 0.998077, \\
V_1/CF_0 &= 9.898, \\
V_2/CF_0 &= 67.030, \quad V/CF_0 = 76.928, \\
D_2^{\text{mod}} &= 9.615 + 66.667 = 76.282, \quad D_1^{\text{mod}} \approx 5.29, \\
w_1 &= 0.1287, \quad w_2 = 0.8713, \\
D_{\text{growth},W}^{\text{mod}} &= 67.145.
\end{aligned}$$

Mature firm at $g = 0.005$: $V/CF_0 = 28.714$, $D^{\text{mod}} = 28.571$. Differencing: $\Delta D_W^{\text{mod}} = 38.574$, $|\Delta \log \text{BM}^W| = 0.986$, so $K_W^C = 39.1$.

A.7 Sign checks

The following internal consistency checks are maintained throughout.

1. High $\log \text{BM}$ implies low $D_{i,G}^{\text{mod}}$: high-BM firms are short-duration.
2. Value sort (long-leg BM > short-leg BM): $D_L^{\text{mod}} < D_S^{\text{mod}}$, so $\beta_k^r > 0$.
3. Positive BM-spread times positive scaling constant: predicted rate-beta positive.
4. Calibration Result 1 calibrations: $K_W > 0$ in all three, preserving the sign.
5. IBES BM-decile spread: D1 ($g_L = 19.78\%$) exceeds D10 ($g_L = 13.17\%$), so the sign of Δg_L matches the two-firm setup.

A.8 Proposition 4 Monte Carlo verification

The construction $X = y^2 + \theta u$ with $y, u \sim N(0, 1)$ independent and $\theta = 2$ satisfies the claims of Proposition 4 by the symmetry and continuity arguments in the proof. Monte Carlo simulation at $N = 2 \times 10^6$ confirms the analytical result: $\text{BM-spread}_k = 0$ to within Monte Carlo noise, and $\Delta u_k \approx 2.87$, so $\beta_k^r \approx K_G \cdot 2.87 \approx 95.7$ under Gordon scaling while the naive spanning prediction is exactly zero.

A.9 Bootstrap confidence interval for the main slope

Resampling signals with replacement and re-running the cross-sectional regression 5,000 times delivers mean slope +1.54 with 95% percentile interval $[+1.06, +2.47]$ and probability of positive slope equal to one. The bootstrap distribution is entirely positive.

A.10 First-stage R^2 for the orthogonalized FFR

Regressing ΔFEDFUNDS on ΔVIX and ΔCredit over 1990–2021 gives $R^2 = 0.061$. Adding ΔTIPS over the 2003–2021 sub-sample gives $R^2 = 0.133$. The correlation between raw ΔFEDFUNDS and the orthogonalized residual (notips version) is 0.969. Orthogonalization removes a small fraction of the variance and preserves most of the shock.