Why Overpricing Persists: Asymmetric Information Quality
and the Cross-Section of Anomaly Correction*

March 23, 2026

Abstract

Analyst career incentives suppress negative research, creating a precision gap between signals about overpriced and underpriced stocks. A model with Bayesian arbitrageurs who receive asymmetric-quality signals generates short-leg dominance: overpricing persists longer than underpricing because arbitrageurs trade less aggressively when their information is poor. In continuous time, individual correction times factor as the product of a type-dependent deterministic half-life and a common random variable, implying that the mean correction-time ratio equals the coverage ratio k, the variance ratio equals k^2, and the coefficient of variation is identical across types. A two-channel decomposition separates information and cost contributions. Joint calibration to momentum portfolios attributes 65% of the total short-leg dominance to the information channel and 35% to costs, with neither channel alone matching all empirical moments. A separating test rejects the prediction that publication compresses the short leg more than the long leg, revealing that the cost channel drives post-publication

*For helpful comments, I thank seminar participants. All errors are my own.

dynamics while the information channel determines the persistent component that survives market development. Empirical tests confirm short-leg dominance for three of four major anomalies, a convex IVOL-alpha relationship, and asymmetric sentiment sensitivity concentrated on the short leg.

JEL Classification: G12, G14, G24

Keywords: short-leg dominance, anomalies, analyst coverage, Bayesian learning, information asymmetry

Contents

1 Introduction

Why do anomaly returns concentrate in the short leg? Across virtually every documented cross-sectional anomaly, the profits from shorting overpriced stocks exceed the profits from buying underpriced stocks (Stambaugh et al., 2012). Momentum losers generate larger absolute alphas than winners. Stocks with weak profitability or aggressive investment underperform more than their counterparts outperform. The short-leg share of anomaly profits averages roughly two-thirds. The standard explanation appeals to short-sale costs: trading against overpricing is expensive. But short-sale costs explain why arbitrageurs trade less, not why they *know* less. Overpricing could persist even in a world with zero short-sale costs if the information available about overpriced stocks is systematically worse than the information available about underpriced stocks.

Analyst career incentives suppress negative research (Hong and Kubik, 2003), creating a precision gap between signals about overpriced and underpriced stocks. Rational arbitrageurs who receive lower-quality signals about overpriced stocks take smaller positions, generate less price impact, and allow overpricing to persist longer. The mechanism operates through information production, not through trading costs, and a formal two-channel decomposition separates the two sources of short-leg dominance. Joint calibration to momentum portfolios attributes 65% of the total asymmetry to the information channel and 35% to costs, with neither channel alone matching all empirical moments.

The model has four parts. Part I establishes a discrete-time benchmark. Diagnostic investors create mispricing in both directions. Rational CARA arbitrageurs receive public signals whose precision depends on the stock's type: overpriced stocks have signal precision τ_O, underpriced stocks have precision $\tau_U > \tau_O$. The coverage ratio $k = \tau_U/\tau_O > 1$ captures the bias in analyst coverage quality. In equilibrium, the short-leg dominance ratio $R(t) = (\pi_0 + t\tau_U)/(\pi_0 + t\tau_O)$ exceeds one, increases over time, and is bounded by k. The short-leg premium (the absolute difference between overpricing and underpricing) is hump-shaped in holding period, peaking at $t^* = \pi_0/\sqrt{\tau_O \tau_U}$.

Part II moves to continuous time to characterize stochastic correction dynamics. For a specific overpriced stock, the correction time is a first-passage time of the posterior mean to a deterministic boundary. Using the Dambis-Dubins-Schwarz theorem and a Brownian bridge representation, the correction time factors as

$$\tau_d = H_d \cdot g(U), \tag{1}$$

where $H_d = \pi_0/\tau_d$ is the deterministic half-life from Part I and U is a random variable whose distribution depends on the initial surplus but not on the stock's type. The factorization implies that mean correction times scale by k, variances scale by k^2, and the coefficient of variation is identical for both legs.

Part III solves a risk-constrained fund's allocation problem across the two legs. The Sharpe ratio of the short leg falls below the Sharpe ratio of the long leg at short horizons, because the signal quality disadvantage dominates. At long horizons, the sheer magnitude of overpricing overcomes the precision disadvantage. The crossover occurs at the same threshold $R(t) = \sqrt{k}$ that maximizes the short-leg premium, because the dominance ratio $R(t)$ is the sufficient statistic for relative attractiveness.

Part IV introduces short-sale costs alongside the information channel and derives a formal two-channel decomposition. The total short-leg dominance ratio decomposes additively into an information component (the precision-driven ratio from Part I) and a cost component. The information share is decreasing over time: at short horizons, information asymmetry dominates; at long horizons, the cost channel becomes relatively more important.

Two calibrations bracket the information channel's contribution. A pure-information calibration to momentum decile portfolios from the Ken French library pins the coverage ratio at $k = 7.04$, matching the observed short-to-long alpha ratio of 2.01 at a one-year horizon but missing non-targeted moments (peak SLP timing, underpricing half-life) by factors of two to three. A joint calibration that targets both the alpha ratio and the overpricing half-life yields

$k = 1.93$ and a short-sale cost c_S that together match all moments, with the information channel accounting for 65% of the total dominance ratio and the cost channel accounting for 35%. The pure-information $k = 7.04$ is consistent with the approximately 7:1 ratio of buy-to-sell recommendations documented by Barber et al. (2006); the joint-calibration $k = 1.93$ implies that the effective precision ratio is roughly 2:1, with the remainder of the recommendation asymmetry reflecting quantity of coverage rather than quality.

Empirical tests using portfolio-level data support the model's core predictions. Short-leg dominance holds for momentum (alpha ratio 2.01), profitability (3.02), and investment (1.87), though not for the value anomaly (0.30). The IVOL-alpha relationship is convex, with the gradient 4.6 times steeper for small stocks than large stocks ($t = -6.71$), consistent with the prediction that information asymmetry amplifies the IVOL effect where analyst coverage is weakest. Sentiment, proxied by Michigan Consumer Sentiment, affects short-leg alphas but not long-leg alphas, with the investment anomaly short-leg sentiment coefficient reaching statistical significance ($t = -2.36$).

A separating test examines the model's prediction that publication should compress short-leg alphas more than long-leg alphas (because publication increases τ_O). The data reject this prediction: post-publication, the long leg of momentum compressed by 63% while the short leg compressed by only 12%. The short-leg share of momentum profits rose from 0.61 pre-publication to 0.79 post-publication, and reaches 0.87 in the most recent decade. The rejection is informative. The information channel explains the steady-state pattern (why overpricing is more persistent than underpricing) while the cost channel explains the post-publication dynamics (why arbitrage capital flows primarily to the long side). The increasing short-leg share over six decades of market development is itself evidence for a friction beyond costs: as trading costs have fallen, the long leg has been arbitraged away, but the short leg persists, consistent with the structural career-incentive friction that the model identifies.

The paper relates to three strands of the literature. First, the limits-to-arbitrage literature (Shleifer and Vishny, 1997; Gromb and Vayanos, 2002; Brunnermeier and Pedersen,

2009) explains why sophisticated investors fail to eliminate mispricing but takes the mispricing as given. Mispricing here arises endogenously from belief errors, and information quality, not just trading costs, determines the speed of correction. Second, the behavioral asset pricing literature (DeLong et al., 1990; Bordalo et al., 2019, 2025) specifies the error process in beliefs but does not model the arbitrageur's problem or the correction dynamics. Embedding diagnostic investors in an equilibrium with rational arbitrageurs delivers a structural mapping from observables to correction speeds. Third, the empirical anomaly literature (Stambaugh et al., 2012, 2015; Stambaugh and Yuan, 2017) documents short-leg dominance and links it to sentiment and arbitrage impediments. Stambaugh et al. (2015) attribute the asymmetry to costs; a complementary information channel is formally separable from costs, and a measurement framework attributes 65% of the total asymmetry to information at a one-year horizon.

The paper also complements Van Binsbergen et al. (2023), who classify anomalies as build-up versus resolution types with distinct dynamics, and Lochstoer and Tetlock (2020), who show that cash flow news rather than discount rate news drives anomaly returns. The information channel proposed here is consistent with both findings: analyst coverage bias affects the speed at which cash flow information is incorporated into prices, and the resulting dynamics differ across anomaly types to the extent that coverage quality varies.

Section 2 presents the model. Section 3 states the main results with economic interpretation and reports calibration and empirical evidence. Section 4 addresses limitations, compares the model to alternatives, and develops policy implications. Section 5 concludes.

2 Model

2.1 Environment and agents

Time. Periods $t = 0, 1, \ldots, T$. At T, fundamental values are revealed.

Assets. A continuum of stocks indexed by $i \in [0, 1]$. Stock i has fundamental value $v_i \sim N(\mu, \sigma_v^2)$, drawn independently across stocks. The prior precision is $\pi_0 = 1/\sigma_v^2$. A risk-free asset yields gross return $R_f = 1$.

Agents. Two types of investors participate in the market:

1. **Diagnostic investors** (mass $\lambda \in (0, 1)$): demand for stock i is $x_i^D = \alpha_D \cdot b_i$, where $b_i = +b > 0$ for overpriced stocks and $b_i = -b$ for underpriced stocks. Demand is inelastic, following the noise trader tradition of DeLong et al. (1990). The parameter b captures the magnitude of diagnostic bias (Bordalo et al., 2019).

2. **Rational arbitrageurs** (mass $1-\lambda$): CARA utility with absolute risk aversion $\gamma > 0$. Each arbitrageur is infinitesimal and takes prices as given.

Assumption 1 (Regularity). $\lambda \alpha_D b < 1$*, so diagnostic demand alone does not exhaust unit supply.*

Assumption 2 (Known mispricing direction). *Arbitrageurs observe the diagnostic demand direction* b_i *for each stock: they know whether a stock faces overpricing pressure (*$b_i > 0$*) or underpricing pressure (*$b_i < 0$*).*

Signal structure. At each period $t \in \{1, \ldots, T\}$, a public signal arrives:

$$y_{it} = v_i + \varepsilon_{it}, \qquad \varepsilon_{it} \sim N(0, 1/\tau_d), \tag{2}$$

where $d \in \{O, U\}$ denotes the stock's type (overpriced or underpriced). Noise terms are independent across periods and stocks. After t signals, cumulative precision is $t \cdot \tau_d$.

2.2 Asymmetric analyst coverage: the micro-foundation

The signal precision τ_d differs between overpriced and underpriced stocks. The source of this difference is the well-documented bias in analyst coverage driven by career incentives.

Analysts face career penalties for negative assessments. Hong and Kubik (2003) show that analysts who issue pessimistic forecasts relative to consensus are more likely to be fired or demoted, while optimistic analysts face smaller penalties. For underpriced stocks, an analyst who identifies the mispricing can issue a buy recommendation, which is career-safe and profitable. Many analysts produce informative research, yielding high coverage quality. For overpriced stocks, an analyst who identifies the overpricing must issue a sell recommendation, which is career-dangerous. Fewer analysts produce genuinely informative research, yielding lower coverage quality.

The number of informative analysts covering a stock of type d is n_d. Baseline coverage is n_0 for all stocks. Underpriced stocks receive $\Delta_n > 0$ additional informative analysts (those willing to issue buy recommendations). Each analyst contributes precision $\rho > 0$:

$$\tau_O = n_0 \rho, \qquad \tau_U = (n_0 + \Delta_n)\rho. \tag{3}$$

Assumption 3 (Coverage bias)**.** *$\Delta_n > 0$, so $\tau_O < \tau_U$. The coverage ratio is*

$$k \equiv \frac{\tau_U}{\tau_O} = 1 + \frac{\Delta_n}{n_0} > 1. \tag{4}$$

The ratio k is increasing in Δ_n (larger career penalties widen the bias) and decreasing in n_0 (high baseline coverage dilutes the marginal bias). Short-sale costs make it expensive to *trade on* information about overpricing; analyst coverage bias makes it harder to *produce* information about overpricing. An arbitrageur who can short for free still faces lower signal quality if the analyst ecosystem does not produce informative sell-side research.

2.3 Arbitrageur's problem

At period t, after observing signals $y_{i1}, \ldots, y_{it}$, the arbitrageur forms a Gaussian posterior:

$$v_i \mid \mathcal{F}_t \sim N\!\left(\hat{\mu}_i(t),\ \frac{1}{\pi_0 + t\tau_d}\right), \tag{5}$$

where $\hat{\mu}_i(t) = (\pi_0 \mu + t\tau_d \bar{y}_i(t))/(\pi_0 + t\tau_d)$ and $\bar{y}_i(t) = t^{-1} \sum_{s=1}^{t} y_{is}$.

The arbitrageur chooses position x_{it}^A to maximize CARA expected utility of terminal wealth $W = W_0 + x(v_i - P_{it})$. Under CARA utility $u(W) = -\exp(-\gamma W)$ and the Gaussian posterior $v_i \mid \mathcal{F}_t \sim N(\hat{\mu}_i(t), 1/(\pi_0 + t\tau_d))$, the expected utility is

$$\mathbb{E}_t\Big[-e^{-\gamma[W_0 + x(v_i - P_{it})]}\Big] = -\exp\Big(-\gamma\Big[W_0 + x(\hat{\mu}_i(t) - P_{it}) - \frac{\gamma x^2}{2(\pi_0 + t\tau_d)}\Big]\Big).$$

Maximizing the certainty equivalent $x(\hat{\mu}_i(t) - P_{it}) - \gamma x^2/[2(\pi_0 + t\tau_d)]$ with respect to x yields the first-order condition:

$$x_{it}^A = \frac{(\hat{\mu}_i(t) - P_{it})(\pi_0 + t\tau_d)}{\gamma}. \tag{6}$$

Demand is proportional to the expected mispricing $\hat{\mu}_i(t) - P_{it}$ and to the posterior precision $\pi_0 + t\tau_d$. Higher precision increases demand both because the expected profit estimate is more reliable and because the position variance is smaller per unit of exposure.

2.4 Equilibrium

Market clearing requires $\lambda \alpha_D b_i + (1 - \lambda) x_{it}^A = 1$. Substituting (6) and solving for the price:

$$P_{it} = \mathbb{E}_t[v_i] - \frac{\gamma(1 - \lambda \alpha_D b_i)}{(1 - \lambda)(\pi_0 + t\tau_d)}. \tag{7}$$

The expected mispricing, defined as the deviation of the expected price under diagnostic demand from the no-bias benchmark, is:

$$m_d(t) = \frac{C}{\pi_0 + t\tau_d}, \qquad C \equiv \frac{\gamma \lambda \alpha_D b}{1 - \lambda}. \tag{8}$$

For overpriced stocks ($b_i = +b$), $m^O(t) = C/(\pi_0 + t\tau_O) > 0$: the price exceeds fundamental value. For underpriced stocks ($b_i = -b$), $|m^U(t)| = C/(\pi_0 + t\tau_U) > 0$: the price falls short of fundamental value.

2.5 Continuous-time extension

The discrete benchmark describes average behavior. To characterize the stochastic properties of individual correction times, the model extends to continuous time.

Signal process. For stock i of type d:

$$dY_i(t) = v_i\,dt + \frac{1}{\sqrt{\tau_d}} dB_i(t), \tag{9}$$

where B_i is a standard Brownian motion. Over the interval $[0, t]$, cumulative signal precision is $t\tau_d$, matching the discrete model.

Filtering. The arbitrageur's posterior mean evolves according to the Kalman-Bucy filter:

$$d\hat{\mu}_i = \frac{\sqrt{\tau_d}}{\pi_d(t)}\,dI_i(t), \tag{10}$$

where $\pi_d(t) = \pi_0 + \tau_d t$ is the posterior precision and $dI_i(t) = \sqrt{\tau_d}[dY_i - \hat{\mu}_i\,dt]$ is the innovation process (a standard Brownian motion under the arbitrageur's filtration).

Individual mispricing. For a specific stock i of type d, define the overestimation:

$$Z_i(t) = \hat{\mu}_i(t) - v_i. \tag{11}$$

The process Z_i is a continuous martingale with time-varying volatility $\sqrt{\tau_d}/\pi_d(t)$. Correction occurs when $Z_i(t)$ falls to the risk-premium boundary $\beta_d(t) = C/\pi_d(t)$.

Assumption 4 (Genuine overpricing)**.** *The initial overestimation exceeds the equilibrium risk premium:* $a = \mu - v_i > C/\pi_0$*. The initial surplus is* $x_0 = a - C/\pi_0 > 0$*.*

Assumption 4 restricts the analysis to stocks with $v_i < \mu - C/\pi_0$, so that the mispricing exceeds the rational risk premium. Stocks with $v_i \in (\mu - C/\pi_0, \mu)$ are technically overpriced but priced within the rational risk premium.

2.6 Short-sale costs: the second channel

To decompose short-leg dominance into information and cost components, a per-unit short-sale cost enters the arbitrageur's payoff.

Assumption 5 (Short-sale costs)**.** *Shorting stock i incurs a per-unit cost $c_S \geq 0$, paid by the arbitrageur. Long positions incur no additional cost.*

Short-sale costs reduce the arbitrageur's shorting demand: for overpriced stocks, the effective payoff becomes $x(v_i - P_i + c_S)$, reducing the optimal short position by $c_S(\pi_0 + t\tau_O)/\gamma$. The equilibrium price of overpriced stocks shifts up by c_S. The modified expected overpricing is $m^O(t) + c_S$.

3 Results

3.1 Part I: Short-leg dominance and correction dynamics

Proposition 1 (Short-leg dominance)**.** *For any $t \geq 1$:*

(a) ***Dominance ratio.*** *$R(t) \equiv m^O(t)/|m^U(t)| = (\pi_0 + t\tau_U)/(\pi_0 + t\tau_O) > 1$, strictly increasing in t, with $R(0) = 1$ and $\lim_{t\to\infty} R(t) = k$.*

(b) ***Half-lives.*** *The number of signal periods to halve mispricing is $H_d = \pi_0/\tau_d$. The decay is hyperbolic ($m_d(t) \propto 1/(\pi_0 + t\tau_d)$, not exponential), so "half-life" refers to the time for the denominator to double, which halves the mispricing level.*[1] *The ratio satisfies $H_O/H_U = k$.*

(c) ***IVOL amplification.*** *$\partial R/\partial(\sigma_v^2) > 0$. High-IVOL stocks exhibit greater short-leg dominance.*

[1] The terminology differs from exponential half-life ($e^{-\lambda t}$). Here, the posterior precision grows linearly, producing $1/t$ rather than e^{-t} decay. At time H_d, the precision has doubled from π_0 to $2\pi_0$, and the mispricing has halved.

All proofs appear in the Appendix.

Lower signal precision ($\tau_O < \tau_U$) yields more diffuse posteriors for overpriced stocks, which yield smaller CARA-optimal short positions (equation 6), which yield less price correction, which yields more persistent overpricing. Each additional signal period widens the cumulative precision gap by $\tau_U - \tau_O$, so the asymmetry compounds over time. High-IVOL stocks have diffuse priors (small π_0), making the signal the primary information source and allowing the full precision asymmetry to manifest.

Proposition 2 (Hump-shaped short-leg premium)**.** *The short-leg premium* $\text{SLP}(t) = m^O(t) - |m^U(t)|$ *satisfies:*

(a) $\text{SLP}(0) = 0$ *and* $\text{SLP}(t) > 0$ *for* $t \geq 1$*.*

(b) $\text{SLP}(t)$ *peaks at* $t^* = \pi_0/\sqrt{\tau_O \tau_U}$*.*

(c) The peak value is $C(\sqrt{k} - 1)/[\pi_0(\sqrt{k} + 1)]$*.*

The hump shape reflects two competing forces. At short horizons, the precision gap has not yet accumulated enough to generate a large absolute difference. At long horizons, both mispricings have largely corrected, so the difference shrinks. The peak occurs when underpricing is substantially corrected but overpricing remains. Cost-based models predict monotone divergence in the short-leg premium; the information channel generates a hump shape.

3.2 Part II: Stochastic correction-time distribution

Proposition 3 (Correction-time factorization)**.** *The calendar-time correction time for a stock of type d factors as*

$$\tau_d = H_d \cdot g(U), \tag{12}$$

where $g(u) = u/(1-u)$*,* $H_d = \pi_0/\tau_d$ *is the deterministic half-life, and* U *is a random variable on* $[0, 1]$ *whose distribution depends on* $\zeta = x_0/\sigma_v$ *but not on* τ_d*.*

The factorization arises from a time-change argument. The overestimation $Z_i(t) = \hat{\mu}_i(t) - v_i$ is a continuous martingale with total quadratic variation σ_v^2. The Dambis-Dubins-Schwarz theorem converts it to a Brownian motion in a transformed time scale $s_d(t) = \sigma_v^2 - 1/\pi_d(t)$. Since $Z_i(0) = a$ and $Z_i(\infty) = 0$, the Brownian motion is conditioned to end at $-a$, producing a Brownian bridge. The correction boundary, expressed in the transformed time, is linear, reducing the correction problem to the first crossing of a standard Brownian bridge through a linear boundary. The crossing time $U \in [0, 1]$ depends on the normalized surplus ζ but not on the signal precision τ_d, which enters only through the deterministic time change. The Appendix provides the full derivation.

Proposition 4 (Asymmetric correction-time moments)**.** *Conditional on correction occurring ($U < 1$):*

(a) $\mathbb{E}[\tau_O]/\mathbb{E}[\tau_U] = k$*.*

(b) $\mathrm{Var}[\tau_O]/\mathrm{Var}[\tau_U] = k^2$*.*

(c) $\mathrm{CV}[\tau_O] = \mathrm{CV}[\tau_U]$ *(coefficient-of-variation invariance).*

(d) For any threshold $\bar{T} > 0$*:* $P(\tau_O > \bar{T}) > P(\tau_U > \bar{T})$ *(right-tail dominance).*

Parts (a)–(c) follow from the factorization: since $\tau_d = H_d \cdot g(U)$ and U does not depend on d, all moments scale by powers of H_d, and scale-free properties are type-independent. The precision asymmetry operates as a pure timescale effect: it stretches all moments of the correction-time distribution by powers of k without changing the distributional shape.

3.3 Part III: Risk-budget allocation

A risk-constrained fund allocates a fixed risk budget K between the long leg and the short leg at each instant t, maximizing expected alpha subject to a variance constraint. The expected alpha for leg d at period t is $\alpha_d(t) = m_d(t) = C/\pi_d(t)$. Each signal period, the return on a unit position in leg d has standard deviation $1/\sqrt{\tau_d}$ (the noise per signal draw). The instantaneous Sharpe ratio per leg is therefore:

$$\mathrm{SR}_d(t) = \frac{\alpha_d(t)}{1/\sqrt{\tau_d}} = \frac{C\sqrt{\tau_d}}{\pi_d(t)}. \tag{13}$$

The Sharpe ratio declines in t for both legs as the alpha $C/\pi_d(t)$ shrinks with accumulating precision, while the per-period return volatility $1/\sqrt{\tau_d}$ is constant. The ratio of Sharpe ratios is:

$$\frac{\mathrm{SR}_O(t)}{\mathrm{SR}_U(t)} = \frac{\sqrt{\tau_O}}{\sqrt{\tau_U}} \cdot \frac{\pi_U(t)}{\pi_O(t)} = \frac{R(t)}{\sqrt{k}}.$$

The short leg has a lower Sharpe ratio when $R(t) < \sqrt{k}$ (short horizons) and a higher Sharpe ratio when $R(t) > \sqrt{k}$ (long horizons).

Proposition 5 (Sharpe ratio crossover)**.** *(a)* $\mathrm{SR}_O(t) < \mathrm{SR}_U(t)$ *for* $t < t^*$*,* $\mathrm{SR}_O(t^*) = \mathrm{SR}_U(t^*)$*, and* $\mathrm{SR}_O(t) > \mathrm{SR}_U(t)$ *for* $t > t^*$*, where* t^* *is the SLP peak time from Proposition 2.*

(b) Both the Sharpe ratio crossover and the SLP peak occur at t^* *because both are governed by the condition* $R(t) = \sqrt{k}$*.*

The ratio $\mathrm{SR}_O/\mathrm{SR}_U = R(t)/\sqrt{k}$. At short horizons, $R(t)$ is close to 1, which falls below $\sqrt{k}$, so the long leg dominates on a risk-adjusted basis. At long horizons, $R(t)$ approaches $k > \sqrt{k}$, so the short leg dominates. The crossover occurs when $R(t) = \sqrt{k}$, which is the same condition that maximizes the SLP. The dominance ratio $R(t)$ is the sufficient statistic for relative attractiveness.

Proposition 6 (Optimal risk allocation)**.** *When the risk constraint binds, the risk-budget share allocated to the short leg is:*

$$w_O(t) = \frac{\mathrm{SR}_O^2(t)}{\mathrm{SR}_O^2(t) + \mathrm{SR}_U^2(t)}. \tag{14}$$

The share satisfies $w_O(t) < 1/2$ *for* $t < t^*$*,* $w_O(t^*) = 1/2$*, and* $w_O(t) > 1/2$ *for* $t > t^*$*.*

Proposition 7 (Welfare cost of precision asymmetry)**.** *The constrained fund's expected alpha under asymmetric precision falls below the symmetric benchmark (with $\tau = \sqrt{\tau_O \tau_U}$) for $t < t_w$, where $t_w > t^*$. At t^*, the symmetric benchmark still dominates because $(\sqrt[4]{k} + \sqrt[4]{k}^{-1})/2 > 1$ by the AM-GM inequality.*

The precision asymmetry is costly at short horizons (degrading the short leg's risk-adjusted return) but beneficial at long horizons (creating a large overpricing opportunity). Since most anomaly strategies operate at horizons below t^*, the relevant regime for practitioners is the welfare-loss regime.

3.4 Part IV: Two-channel decomposition

Proposition 8 (Information and cost decomposition)**.** *Under Assumption 5, the total short-leg dominance ratio decomposes as:*

$$R_{TC}(t) = \frac{m^O(t) + c_S}{|m^U(t)|} = R_I(t) + \frac{c_S \cdot \pi_U(t)}{C}, \tag{15}$$

where $R_I(t) = \pi_U(t)/\pi_O(t)$ is the information-only ratio. The information share $S_I(t) = R_I(t)/R_{TC}(t)$ is strictly decreasing in t.

At short horizons (t near 0), information asymmetry accounts for nearly all of the dominance ratio because the cost component $c_S \pi_U(t)/C$ is small. At long horizons, the cost component grows without bound (since $\pi_U(t) = \pi_0 + t\tau_U$ grows linearly) while the information ratio converges to k. The information channel has its greatest relative importance early, before the cumulative cost disadvantage of shorting has fully compounded.

3.5 Calibration

The model has six economically meaningful parameters: the coverage ratio k, the IVOL σ_v, the overpriced signal precision τ_O, the diagnostic investor share λ, the arbitrageur risk aversion γ, and the diagnostic demand scale $\alpha_D b$ (normalized to 1). Five target moments from the data identify the parameters.

Data and targets. Momentum decile portfolios from the Ken French library (July 1963 to December 2024, 738 months) serve as the primary calibration anomaly. CAPM alphas of the short leg (D1, losers) and long leg (D10, winners) yield a long-short alpha of 1.38%/month (16.6% annualized) and a short-to-long absolute alpha ratio of 2.01. Median annualized idiosyncratic volatility is 30% (Ang et al., 2006). McLean and Pontiff (2016) document a 58% post-publication alpha decay. Stambaugh et al. (2012) attribute approximately 60% of anomaly alpha to sentiment-sensitive periods.

Table 1: Calibrated parameters

Parameter	Value	Identification	Source
k (coverage ratio)	7.04	Short/long alpha ratio $R(4) = 2.01$	Internal
σ_v (IVOL)	0.30	Median firm-level IVOL	Ang et al. (2006)
τ_O (overpriced precision)	0.556	Overpricing half-life $H_O = 20$ quarters	Internal
τ_U (underpriced precision)	3.913	$k \times \tau_O$	Derived
λ (diagnostic share)	0.60	Sentiment fraction of alpha	Stambaugh et al. (2012)
γ (risk aversion)	163.1	Level of L-S alpha	Internal

The implied $k = 7.04$ matches the approximately 7:1 ratio of buy-to-sell recommendations reported by Barber et al. (2006).

Model fit: pure-information calibration. Table 2 (Panel A) compares model-implied moments to data targets under the pure-information calibration ($c_S = 0$). The model matches the four targeted moments exactly (the system is exactly identified, with zero overidentifying restrictions). The non-targeted peak SLP timing of 1.9 years falls at the low end of the 3–5 year empirical range, and the implied underpricing half-life of 2.8 quarters is shorter than empirical estimates of roughly 8–10 quarters. Both discrepancies point to the same source: the pure-information model attributes the full short-to-long asymmetry to information quality, while in reality short-sale costs contribute as well.

Joint calibration: information plus costs. The two-channel decomposition (Proposition 8) enables a joint calibration that identifies both channels simultaneously. Adding a fifth target, the overpricing half-life $H_O = 8$ quarters (from event-time anomaly return estimates), the system becomes overidentified. The overpricing half-life pins $\tau_O = \pi_0/H_O = 0.0771$, which with $\tau_U = 0.149$ gives $k = 1.93$. The information-only ratio at $t = 4$ quarters is $R_I(4) = 1.31$, which falls short of the observed 2.01. The short-sale cost absorbs the residual: $c_S = C(2.01 - 1.31)/(\pi_0 + 4\tau_U) = 140.9$ (in model units). The information share at the one-year horizon is $S_I(4) = 1.31/2.01 = 0.65$, or 65%.

Panel B of Table 2 reports the joint-calibration fit. The joint model matches all targeted moments and substantially improves the non-targeted moments: peak SLP timing rises to 3.1 years (within the 3–5 year empirical range), and the underpricing half-life extends to 4.1 quarters (closer to the empirical 8–10 quarters, with the remaining gap attributable to the constant-cost specification).

Table 2: Model fit

Moment	Pure-info	Joint	Data	Targeted?
Panel A: Targeted moments				
L-S CAPM alpha (%/yr)	16.6	16.6	16.6	Yes
|Short|/|Long| alpha ratio	2.01	2.01	2.01	Yes
Half-life ratio H_O/H_U	7.04	–	7.04	Yes (pure-info)
Post-publication decay fraction	0.58	0.58	0.58	Yes
Overpricing half-life H_O (quarters)	20.0	8.0	$\sim$8	Yes (joint)
Panel B: Non-targeted moments				
Peak SLP timing (years)	1.9	3.1	3–5	No
H_U (quarters)	2.8	4.1	$\sim$8–10	No
Panel C: Key parameters				
k (coverage ratio)	7.04	1.93	–	–
c_S (short-sale cost)	0	140.9	–	–
Information share S_I at 1 year	100%	65%	–	–

Two results emerge from the joint calibration. First, neither channel alone suffices: the pure-information model misses the non-targeted moments by factors of two to three, and a pure-cost model ($k = 1$) produces a counterfactual time-series pattern (the dominance ratio grows without bound rather than converging). Second, the information channel accounts for roughly two-thirds of the total asymmetry at a one-year horizon, with the cost channel contributing the remaining third. The joint-calibration $k = 1.93$ implies that the effective precision of signals about underpriced stocks is roughly twice that of overpriced stocks, corresponding to a standard deviation of analyst forecast errors for overpriced stocks that is $\sqrt{k} \approx 1.4$ times larger than for underpriced stocks. The information share $S_I(t)$ declines with horizon: 65% at 1 year, 58% at 2 years, 49% at 5 years, and converging to $k/(k + c_S \pi_O(t)/C) \to 0$ as the cost channel dominates at long horizons. The information channel is most important at short to medium horizons where signal quality matters most; the cost channel increasingly binds at longer horizons as the cumulative cost of maintaining short positions grows. The pure-information $k = 7.04$ matches the approximately 7:1 buy-to-sell recommendation ratio (Barber et al., 2006), suggesting that the recommendation count ratio is an upper bound on the precision ratio (the number of bullish analysts exceeds the number producing genuinely informative research on either side).

Sensitivity. Table 3 reports the effect of 20% perturbations in each parameter. The IVOL parameter σ_v is the most influential: a 20% increase nearly doubles the long-short alpha. The diagnostic share λ scales the mispricing level but does not affect the alpha ratio or timing. The separation between level parameters (λ, γ) and asymmetry parameters (k, σ_v) follows from the linearity of diagnostic demand in the model. With price-elastic diagnostic demand, the separation would break because the diagnostic investors' response would depend on the equilibrium price, which in turn depends on k and σ_v.

3.6 Empirical evidence

Five tests use portfolio-level data from the Ken French library and FRED (July 1963 to December 2024). Table 4 summarizes the results.

Table 3: Sensitivity analysis

Parameter	Perturbation	L-S alpha (%/yr)	Alpha ratio	Peak SLP (yr)
Baseline	–	16.6	2.01	1.9
k	+20%	17.5	2.24	1.7
k	−20%	15.3	1.77	2.1
σ_v	+20%	28.3	2.35	1.3
σ_v	−20%	8.3	1.69	2.9
λ	±20%	16.6	2.01	1.9

Table 4: Empirical tests: summary

Test	Result	Support
Short-leg dominance	Ratio 2.01 (momentum), 3.02 (profitability), 1.87 (investment), 0.30 (value)	3 of 4
Sentiment asymmetry	Short-leg alpha 38–48 bps/mo more negative in high-sentiment months; long-leg unchanged	Partial
IVOL gradient	Small-stock gradient 4.6× big-stock; $t = -6.71$	Yes
Horizon profile	Short-leg share slightly decreases 0.67 to 0.63	No (weak test)
IVOL-alpha convexity	Negative quadratic coefficient, $p = 0.028$	Yes

Short-leg dominance across anomalies. CAPM alphas of extreme decile portfolios confirm short-leg dominance for momentum (−0.92%/month vs. +0.46%/month, ratio = 2.01), profitability (−0.36% vs. +0.12%, ratio = 3.02), and investment (−0.29% vs. +0.16%, ratio = 1.87). The value anomaly is the exception: the growth-stock decile has a near-zero alpha (−0.06%, $t = -0.74$), and the value premium is long-leg driven. Section 4 discusses why the information mechanism may not apply to value.

IVOL gradient asymmetry. Using 25 portfolios double-sorted on size and residual variance, the IVOL-alpha gradient (Hi RESVAR minus Lo RESVAR alpha) is −1.58%/month for small stocks versus −0.34%/month for large stocks, a ratio of 4.64 ($t = -6.71$ on the difference). The model predicts the gradient ratio equals $R(t)^2$, and $R(4)^2 = 2.01^2 = 4.04$, close to the observed 4.64.

IVOL-alpha convexity. Across ten residual variance deciles, the alpha profile is flat through D8 and then drops sharply: D10 has an alpha of -0.90%/month while D1 through D8 are all near zero. A quadratic fit yields a significantly negative quadratic coefficient ($F = 7.68$, $p = 0.028$). The convex functional form is consistent with the model's $1/(\pi_0 + t\tau_d)$ mispricing, where low π_0 (high IVOL) generates nonlinear amplification. Linear cost-based models predict a more uniform gradient.

Sentiment sensitivity. Splitting the sample by median Michigan Consumer Sentiment, short-leg alphas are 38–48 basis points per month more negative in high-sentiment months, while long-leg alphas are essentially unchanged. The investment anomaly short-leg coefficient is statistically significant ($t = -2.36$); profitability is marginally significant ($t = -1.74$); momentum has the correct sign but is insignificant ($t = -0.90$). The asymmetric pattern is consistent with the model: diagnostic demand (proxied by sentiment) interacts with the precision gap to generate larger overpricing when sentiment is high.

3.7 Separating test: post-publication alpha compression

The tests above are consistent with the information channel but do not discriminate against cost-based alternatives. A separating test exploits the model's prediction about the dynamics of alpha compression after anomaly publication.

Prediction. Publication increases analyst coverage of previously overlooked overpriced stocks, raising τ_O. Because the marginal correction from increasing τ_O exceeds that from increasing τ_U (the short leg has more room to correct), the model predicts that post-publication alpha compression concentrates on the short leg. Cost-based models predict the opposite: publication attracts arbitrage capital, and long-side arbitrage (buying underpriced stocks) scales more easily than short-side arbitrage, so the long leg should compress more.

Test design. Using momentum decile portfolios from the Ken French library (July 1963 to December 2024), CAPM alphas for the short leg (D1, losers) and long leg (D10, winners) are computed in sub-periods defined by the publication of Jegadeesh and Titman (1993): pre-publication (1963–1992) and post-publication (1994–2024). An interaction regression captures the change in alpha after publication, and a block bootstrap (block size = 12 months, 5,000 replications) provides a confidence interval for the compression difference.

Results. Table 5 reports the sub-period alphas. The long leg compressed by 63% (from +0.66%/month to +0.25%/month, $t = -1.85$ on the shift), while the short leg compressed by only 12% (from −1.03%/month to −0.90%/month, $t = 0.37$ on the shift). The short-leg share of total anomaly profits rose from 0.61 pre-publication to 0.79 post-publication, reaching 0.87 in the most recent decade (2010–2024).

Table 5: Post-publication alpha compression: momentum

Period	Short alpha (%/mo)	Long alpha (%/mo)	|S|/|L|	Short share
Pre-publication (1963–1992)	−1.03	+0.66	1.55	0.61
Post-publication (1994–2024)	−0.90	+0.25	3.67	0.79
Post-2000 (2001–2024)	−0.63	+0.24	2.60	0.72
Recent (2010–2024)	−1.01	+0.14	7.33	0.88
Compression (% of pre-pub)	12%	63%		
Bootstrap p(short > long, one-sided)		0.216		

Interpretation. The data reject the model's prediction that publication compresses the short leg more than the long leg. Post-publication arbitrage capital flows primarily to the long side, consistent with cost-based models where buying is cheaper than shorting.

The rejection is informative about the relative roles of the two channels. The information channel explains the steady-state pattern: why overpricing is more persistent than underpricing in any given regime. The cost channel explains the post-publication dynamics: when arbitrage capital enters following publication, it flows to the long side because long-side trading faces lower costs. The increasing short-leg share over six decades of market development (from 0.50 in the 1960s to 0.87 in the 2020s) points to a friction beyond costs. As trading costs have fallen through decimalization, electronic trading, and ETF creation, the long leg has been progressively arbitraged away. The short leg persists, consistent with the structural career-incentive friction that is resistant to improvements in market infrastructure.

The time-trend in the short-leg share also illustrates an important asymmetry in how the two channels respond to market development. Secular improvements in trading technology reduce c_S, compressing the cost component of short-leg dominance. But the career incentive that suppresses negative analyst research is a labor-market friction, not a trading friction, and may be unaffected by (or even strengthened by) technological improvements in trading. The information channel and the cost channel are complementary, not competing, explanations of short-leg dominance. The information channel accounts for the persistent component that survives market development; the cost channel accounts for the component that erodes over time.

4 Discussion

4.1 Relationship to existing models

Several existing frameworks emerge as special cases.

Symmetric precision recovers the standard noise-trader model. When $k = 1$ ($\tau_O = \tau_U$), the dominance ratio equals one at all horizons, half-lives are identical, and the risk-budget allocation splits equally between legs. Short-leg dominance vanishes. The model reduces to a standard Bayesian learning environment with noise traders (DeLong et al., 1990).

Cost-based models generate short-leg dominance through a different channel. Stambaugh et al. (2015) attribute short-leg dominance to the costs of shorting: when $k = 1$

and $c_S > 0$, the total dominance ratio $R_{TC}(t) = 1 + c_S \pi_U(t)/C$ exceeds one and grows linearly without bound. The information model generates a bounded, concave dominance ratio. The two channels are additive (Proposition 8) and, in principle, separately identifiable from the horizon profile of $R(t)$: the information channel contributes a bounded, concave component; the cost channel contributes an unbounded, linear component.

Symmetric signal reduction differs from directional precision asymmetry. Da et al. (2024) model anomaly alpha persistence as a consequence of high-dimensional estimation noise that symmetrically reduces signal precision. Their mechanism explains why anomaly alphas persist overall but cannot generate the directional asymmetry between legs, because symmetric noise yields $R(t) = 1$. The directional complement is here: overpricing correction is slower than underpricing correction, taking the level of signal noise as given.

Coordination-based models differ in cross-sectional predictions. Garleanu et al. (forthcoming) generate persistent overpricing from coordination failure among short sellers. Individual shorts are unprofitable without others participating. Their model predicts that short-interest concentration matters for correction speed. The information model predicts that analyst coverage quality matters. The two models differ in policy implications: the coordination model suggests improving short-interest transparency would accelerate correction, while the information model suggests improving sell-side research incentives would accelerate correction.

4.2 The value anomaly and observable criteria for k

The empirical tests show that the value anomaly does not exhibit short-leg dominance: the growth-stock decile (the short leg) has a near-zero CAPM alpha (-0.06%/month, $t = -0.74$), and the value premium is concentrated in the long leg. $R > 1$ requires $k > 1$, so the value anomaly is consistent with $k \leq 1$ for this specific anomaly, not a failure of the model's logic.

Short-leg dominance requires $k > 1$, which holds when analyst career incentives suppress negative research about the short-leg stocks. An observable proxy for k is the ratio of informative buy recommendations to informative sell recommendations among analysts covering a given set of stocks. For momentum losers, profitability laggards, and aggressive investors, the stocks in the short leg have low analyst coverage, high forecast dispersion, and very few sell recommendations. Analyst recommendation data from I/B/E/S confirm that sell recommendations are rare for these stocks. For growth stocks (the short leg of the value anomaly), by contrast, technology and high-growth firms attract extensive and often skeptical analyst coverage, including prominent sell-side calls against overvalued names. The ratio of buy-to-sell recommendations for growth stocks is closer to 3:1 than to 7:1, and much of the sell-side research is genuinely informative.

A cross-anomaly prediction follows: anomalies whose short-leg stocks have a high buy-to-sell recommendation ratio (i.e., few informative sell recommendations) should exhibit strong short-leg dominance (k large, $R(t) \gg 1$). Anomalies whose short-leg stocks attract balanced or skeptical coverage should exhibit weak short-leg dominance or long-leg dominance ($k \approx 1$ or $k < 1$). For value stocks, long-leg dominance follows because $k < 1$: the long leg (value stocks) consists of unglamorous firms with thin analyst coverage, while the short leg (growth stocks) has dense coverage. The observed alpha ratio of 0.30 (long-leg dominant) is consistent with this prediction.

4.3 Limitations

Several features of the model warrant explicit discussion.

The Brownian bridge approximation requires a long horizon. The continuous-time results (Part II) use a Brownian bridge representation that is exact in the limit $T \to \infty$. At finite horizons, the fraction of total learning completed by T is $\tau_d T/(\pi_0 + \tau_d T)$. For overpriced stocks with the calibrated parameters ($\tau_O = 0.556$, $\pi_0 = 11.11$), this fraction is approximately 50% at $T = 20$ quarters. The bridge approximation is therefore reasonable for underpriced stocks (88% at 5 years) but less precise for overpriced stocks at moderate horizons. The factorization and moment results hold exactly for any correction event that occurs before T; the approximation affects only the tail probabilities.

The distribution of the crossing time U has no closed-form density. Proposition 3 establishes the factorization, and Proposition 4 derives moment ratios from it, but the density of U does not have a closed-form expression. Numerical computation via eigenfunction expansions of the associated Ornstein-Uhlenbeck first-passage problem (Linetsky, 2004) falls outside the scope of this paper. Absent the density, the model cannot generate absolute correction-time predictions (e.g., "overpriced stocks correct in $X \pm Y$ years"). The moment *ratios* across types are sharp, but the moment *levels* require numerical methods.

The factorization $\tau_d = H_d \cdot g(U)$ does, however, provide a natural estimation strategy. Event-time anomaly return data (individual stock returns tracked from portfolio formation without rebalancing) contain information about the correction-time distribution. The factorization implies that the empirical distribution of correction times for overpriced stocks should be a scaled version of the distribution for underpriced stocks, with the scaling factor equal to k. Estimating the two distributions and testing the scaling property provides a direct test of the factorization. The coefficient-of-variation invariance (Proposition 4c) provides an additional overidentifying restriction: the CV should be identical for both legs, regardless of the value of k.

Risk-budget allocation is static. The fund's problem in Part III is solved at each instant t as a static mean-variance optimization. A fully dynamic formulation would have the fund solve an HJB equation, accounting for the evolving opportunity set and the stochastic possibility of correction. A fund that knows the dominance ratio $R(t)$ is increasing would tilt more toward the short leg than the myopic solution prescribes, because the future opportunity set favors the short leg. The static allocation therefore provides a lower bound on the optimal short-leg share at short horizons and an upper bound at long horizons. The qualitative insight (the long leg dominates at short horizons, the short leg at long horizons, with a crossover at $R(t) = \sqrt{k}$) survives dynamic optimization, but the exact crossover time and the welfare cost magnitude depend on the fund's planning horizon and discount rate.

Career incentive asymmetry is modeled in reduced form. Assumption 3 takes $\tau_O < \tau_U$ as a primitive. A micro-foundation from a cheap-talk model with career-concerned analysts (Hong and Kubik, 2003) would endogenize k as a function of the analyst's payoff structure. The reduced-form approach is grounded in robust empirical evidence on the buy-sell recommendation ratio but does not allow the coverage ratio to respond endogenously to market conditions.

The quantitative mapping from career incentives to the precision ratio deserves scrutiny. The approximately 7:1 ratio of buy-to-sell recommendations (Barber et al., 2006) measures the quantity of coverage by direction, not the quality. An analyst who issues a buy recommendation on an underpriced stock may produce a signal of the same precision as an analyst who issues a sell recommendation on an overpriced stock. The recommendation count ratio is therefore an upper bound on the precision ratio k, because it conflates informative and uninformative recommendations. The joint calibration yields $k = 1.93$, which is substantially below 7:1, consistent with this interpretation: most of the recommendation asymmetry reflects quantity (many analysts issue buy recommendations regardless of information content) rather than quality (the precision of the average informative signal). A formal micro-foundation that maps the career concern payoff structure to signal precision would pin down the relationship between the recommendation ratio and k, but is beyond the scope of the present model.

The model applies to information-driven anomalies. The mechanism operates through the precision gap in analyst coverage. For anomalies where the short leg consists of stocks with high analyst coverage (e.g., the value anomaly, where growth stocks are heavily covered), the information channel is weak. The model's predictions apply most directly to momentum, profitability, and investment anomalies, where the short leg consists of stocks with poor or biased coverage.

4.4 Testable predictions and separating tests

Several predictions distinguish the information channel from cost-based alternatives.

Correction-time variance ratio. The variance of overpriced correction times is k^2 times the variance for underpriced stocks (Proposition 4b). Cost-based models make no prediction about the variance ratio. A test requires firm-level correction-time data, which can be constructed from event-time returns tracked without rebalancing.

Coefficient-of-variation invariance. The CV of correction times is identical for both legs (Proposition 4c). If the information mechanism is the primary channel, this invariance should hold empirically. If the cost channel dominates, CV should differ because costs affect only the short leg.

Hump-shaped short-leg premium. The short-leg premium peaks at $t^* \approx 3.1$ years under the joint calibration (1.9 years under the pure-information calibration) and then declines. Cost-based models predict monotone divergence. The horizon profile at 1, 2, 3, 5, and 10 years post-formation discriminates between the two channels. The test requires firm-level event-time returns rather than rebalanced portfolio returns.

Post-publication compression. Publication increases τ_O (more analysts cover previously overlooked overpriced stocks). The marginal correction from increasing τ_O exceeds that from increasing τ_U because $\partial m^O/\partial\tau_O > \partial|m^U|/\partial\tau_U$ when $\tau_O < \tau_U$. The model therefore predicts a larger fraction of alpha decline on the short leg. As documented in Section 3.7, the data reject this prediction for momentum: the long leg compressed far more than the short leg after publication. The rejection suggests that publication operates primarily through the cost channel (attracting arbitrage capital to the long side) rather than through the information channel (improving analyst coverage of overpriced stocks). The career incentive friction that suppresses negative research may be resistant to the dissemination effects of publication.

Within-anomaly cross-sectional tests. Stocks with lower analyst coverage quality within the short leg exhibit more persistent overpricing. Using I/B/E/S data, stocks within the momentum loser decile can be sorted by the number of sell recommendations, forecast dispersion, or the buy-to-sell ratio. Losers with fewer sell recommendations (higher firm-level k) have more persistent negative alphas than losers with balanced coverage. Within-anomaly tests of this type exploit cross-sectional variation in k rather than the cross-anomaly variation in Table 4 and provide stronger identification of the information channel.

MiFID II natural experiment. MiFID II (January 2018) changed research payments in European markets, potentially affecting sell-side research quality. Guo and Mota (2021) document significant declines in analyst coverage for small European firms following MiFID II, and Fang et al. (2020) find that the unbundling of research payments reduced the number of analysts covering affected stocks. MiFID II's effect on $R(t)$ depends on whether the coverage reduction disproportionately affects informative sell-side research (increasing k and worsening short-leg dominance) or reduces coverage symmetrically (leaving k unchanged). A difference-in-differences design comparing European and U.S. stocks before and after January 2018 identifies the net effect.

4.5 Policy implications

The information channel points to policy levers absent from cost-based models.

First, improving sell-side research incentives (e.g., regulatory protection for bearish analysts, buy-side funding for independent research) reduces k and compresses short-leg dominance. Reducing short-sale costs, by contrast, does not eliminate short-leg dominance if the information asymmetry persists.

Second, mandatory disclosure improvements (Reg FD, SOX, IFRS adoption) reduce information asymmetry and should compress the short-leg premium. The marginal social value of transparency is higher for overpriced stocks: an additional informative analyst covering overpriced stocks produces a larger correction than an additional analyst covering underpriced stocks.

Third, the welfare cost of precision asymmetry (Proposition 7) implies that market efficiency is better served by improving information quality on the short side than by reducing trading costs on the short side, at least at horizons below t^*.

5 Conclusion

Analyst career incentives create a precision gap that makes overpricing harder to identify and slower to correct than underpricing. The precision gap generates short-leg dominance, a hump-shaped short-leg premium, and a correction-time factorization into a type-dependent timescale and a common random multiplier. A two-channel decomposition separates information and cost contributions to short-leg dominance. Joint calibration to momentum portfolios attributes 65% of the total asymmetry to the information channel and 35% to costs, with neither channel alone matching all empirical moments. A separating test rejects the prediction that publication compresses the short leg more than the long leg, revealing that the cost channel drives post-publication dynamics while the information channel determines the persistent component that survives market development. The information channel complements rather than replaces cost-based explanations, and points to a distinct policy lever: improving the quality of sell-side research on overpriced stocks addresses the informational root of persistent overpricing that cost reduction alone cannot reach.

A Proofs

A.1 Proof of Proposition 1 (Short-leg dominance)

Part (a). From equation (8):

$$R(t) = \frac{m^O(t)}{|m^U(t)|} = \frac{\pi_0 + t\tau_U}{\pi_0 + t\tau_O}.$$

For $t \geq 1$, $\tau_U > \tau_O$ implies $\pi_0 + t\tau_U > \pi_0 + t\tau_O$, so $R(t) > 1$.

$R(t)$ is strictly increasing:

$$\frac{dR}{dt} = \frac{\pi_0(\tau_U - \tau_O)}{(\pi_0 + t\tau_O)^2} > 0.$$

Limits: $R(0) = 1$ and $\lim_{t\to\infty} R(t) = \tau_U/\tau_O = k$.

Part (b). Setting $m_d(H_d) = m_d(0)/2$:

$$\frac{\pi_0}{\pi_0 + H_d\tau_d} = \frac{1}{2} \implies H_d = \frac{\pi_0}{\tau_d}.$$

The ratio: $H_O/H_U = \tau_U/\tau_O = k$.

Part (c). Write $\pi_0 = 1/\sigma_v^2$. Then:

$$\frac{dR}{d(\sigma_v^2)} = \frac{dR}{d\pi_0} \cdot \frac{d\pi_0}{d(\sigma_v^2)}.$$

$dR/d\pi_0 = t(\tau_O - \tau_U)/(\pi_0 + t\tau_O)^2 < 0$ and $d\pi_0/d(\sigma_v^2) = -1/\sigma_v^4 < 0$. The product is positive.

□

A.2 Proof of Proposition 2 (Hump-shaped short-leg premium)

$$\text{SLP}(t) = m^O(t) - |m^U(t)| = \frac{Ct(\tau_U - \tau_O)}{(\pi_0 + t\tau_O)(\pi_0 + t\tau_U)}.$$

Clearly $\text{SLP}(0) = 0$ and $\text{SLP}(t) > 0$ for $t > 0$ (Part (a)).

For Part (b), differentiate. Let $f(t) = (\pi_0 + t\tau_O)(\pi_0 + t\tau_U)$. Then $\mathrm{SLP}(t) = C(\tau_U - \tau_O) \cdot t/f(t)$. The derivative $d(t/f)/dt = (f - tf')/f^2$. Compute:

$$\begin{aligned} f(t) &= \pi_0^2 + \pi_0 t(\tau_O + \tau_U) + t^2\tau_O\tau_U, \\ f'(t) &= \pi_0(\tau_O + \tau_U) + 2t\tau_O\tau_U, \\ f(t) - tf'(t) &= \pi_0^2 - t^2\tau_O\tau_U. \end{aligned}$$

Setting $f(t) - tf'(t) = 0$: $t^* = \pi_0/\sqrt{\tau_O\tau_U}$.

For $t < t^*$: $\pi_0^2 > t^2\tau_O\tau_U$, so $d\mathrm{SLP}/dt > 0$. For $t > t^*$: $d\mathrm{SLP}/dt < 0$. Since $\mathrm{SLP}(0) = \mathrm{SLP}(\infty) = 0$ and $\mathrm{SLP} > 0$ for $t > 0$, the function is hump-shaped.

For Part (c), substitute t^* into $\mathrm{SLP}(t)$:

$$\begin{aligned} \pi_0 + t^*\tau_O &= \pi_0(1 + 1/\sqrt{k}), \\ \pi_0 + t^*\tau_U &= \pi_0(1 + \sqrt{k}). \end{aligned}$$

Using $\sqrt{k}(1 + 1/\sqrt{k}) = 1 + \sqrt{k}$:

$$\mathrm{SLP}(t^*) = \frac{C(\sqrt{k} - 1)}{\pi_0(\sqrt{k} + 1)}.$$

□

A.3 Proof of Proposition 3 (Correction-time factorization)

Quadratic variation of Z_i. From the filtering equation (10) and $dZ_i = d\hat{\mu}_i$:

$$\langle Z_i \rangle(t) = \int_0^t \frac{\tau_d}{\pi_d(s)^2}\, ds = \frac{1}{\pi_0} - \frac{1}{\pi_d(t)} \equiv s_d(t),$$

with $s_d(t)$ increasing from 0 to σ_v^2 as t goes from 0 to ∞.

Time change. The inverse is $t_d(s) = s\pi_0/[\tau_d(\sigma_v^2 - s)]$ for $s \in [0, \sigma_v^2)$. By the Dambis-Dubins-Schwarz theorem, there exists a Brownian motion W such that $Z_i(t) = a + W(s_d(t))$. Since $Z_i(\infty) = 0$, we have $W(\sigma_v^2) = -a$, so W is a Brownian bridge from 0 to $-a$ over $[0, \sigma_v^2]$.

Boundary in transformed time. Correction occurs at the first time $Z_i(t)$ reaches $C/\pi_d(t)$ from above. In the s-scale, the condition $Z_i(t) = C/\pi_d(t)$ becomes:

$$a(1 - s/\sigma_v^2) + \mathrm{BB}_0(s) = (C/\pi_0)(1 - s/\sigma_v^2),$$

where BB_0 is a standard Brownian bridge (start 0, end 0 at σ_v^2). Rearranging, correction occurs at the first crossing:

$$\mathrm{BB}_0(s) = -x_0(1 - s/\sigma_v^2),$$

where $x_0 = a - C/\pi_0 > 0$. Since $Z_i(0) > C/\pi_0$ (Assumption 4), the process starts above the boundary, and correction is the first time the process hits the boundary from above. In the s-scale, this is the first-passage time $\inf\{s \in [0, \sigma_v^2] : \mathrm{BB}_0(s) \le -x_0(1 - s/\sigma_v^2)\}$.

Normalization. Setting $u = s/\sigma_v^2 \in [0, 1]$ and using $\mathrm{BB}_0(u\sigma_v^2) \overset{d}{=} \sigma_v \cdot \mathrm{BB}^{\mathrm{std}}(u)$:

$$\mathrm{BB}^{\mathrm{std}}(u) \le -\zeta(1 - u), \qquad \zeta = x_0/\sigma_v.$$

The crossing time $U = \inf\{u \in [0, 1] : \mathrm{BB}^{\mathrm{std}}(u) \le -\zeta(1 - u)\}$ depends on ζ but not on τ_d.

Calendar-time factorization. At $s = U\sigma_v^2$:

$$\tau_d = t_d(U\sigma_v^2) = \frac{U\sigma_v^2\pi_0}{\tau_d\sigma_v^2(1 - U)} = \frac{U}{1 - U} \cdot \frac{\pi_0}{\tau_d} = g(U) \cdot H_d.$$

□

A.4 Proof of Proposition 4 (Asymmetric correction-time moments)

Since $\tau_d = H_d \cdot g(U)$ and U does not depend on d:

Part (a). $\mathbb{E}[\tau_d] = H_d \mathbb{E}[g(U)]$, so $\mathbb{E}[\tau_O]/\mathbb{E}[\tau_U] = H_O/H_U = k$.

Part (b). $\mathrm{Var}[\tau_d] = H_d^2 \mathrm{Var}[g(U)]$, so $\mathrm{Var}[\tau_O]/\mathrm{Var}[\tau_U] = H_O^2/H_U^2 = k^2$.

Part (c). $\mathrm{CV}[\tau_d] = \sqrt{\mathrm{Var}[g(U)]}/\mathbb{E}[g(U)]$, independent of d.

Part (d). $P(\tau_d > \bar{T}) = P(g(U) > \bar{T}/H_d) = P(U > \bar{T}/(\bar{T} + H_d))$. Since $H_O > H_U$, $\bar{T}/(\bar{T} + H_O) < \bar{T}/(\bar{T} + H_U)$, and $P(U > \cdot)$ is decreasing. □

Finiteness of moments. The conditional moments $\mathbb{E}[g(U) \mid U < 1]$ and $\mathrm{Var}[g(U) \mid U < 1]$ are finite. Near $u = 1$, the Brownian bridge has standard deviation $\sim \sqrt{1-u}$ while the boundary $-\zeta(1-u)$ falls faster. Using the OU representation $\mathrm{BB}^{\mathrm{std}}(u) = \sqrt{1-u} \cdot X(-\log(1-u))$ where X is an Ornstein-Uhlenbeck process, the first-passage density has exponential tail decay, so $u/(1-u)$ times the density is integrable.

A.5 Proof of Proposition 5 (Sharpe ratio crossover)

Part (a). From (13):

$$\frac{\mathrm{SR}_O}{\mathrm{SR}_U} = \sqrt{\frac{\tau_O}{\tau_U}} \cdot \frac{\pi_U(t)}{\pi_O(t)} = \frac{1}{\sqrt{k}} \cdot R(t).$$

Setting $\mathrm{SR}_O = \mathrm{SR}_U$: $R(t) = \sqrt{k}$.

Solving $(\pi_0 + t\tau_U)/(\pi_0 + t\tau_O) = \sqrt{k}$:

$$\pi_0(1 - \sqrt{k}) = t\tau_O(\sqrt{k} - k) = -t\tau_O\sqrt{k}(\sqrt{k} - 1).$$

Since $\sqrt{k} > 1$: $\pi_0 = t\tau_O\sqrt{k}$, giving $t_{\mathrm{cross}} = \pi_0/(\tau_O\sqrt{k})$.

From Proposition 2: $t^* = \pi_0/\sqrt{\tau_O\tau_U} = \pi_0/(\tau_O\sqrt{k})$. So $t_{\mathrm{cross}} = t^*$.

Part (b). Verify $R(t^*) = \sqrt{k}$:

$$R(t^*) = \frac{\pi_0 + t^*\tau_U}{\pi_0 + t^*\tau_O} = \frac{1+\sqrt{k}}{1+1/\sqrt{k}} = \frac{\sqrt{k}(1+\sqrt{k})}{\sqrt{k}+1} = \sqrt{k}.$$

□

A.6 Proof of Proposition 6 (Optimal risk allocation)

With binding constraint, the Lagrangian yields FOC: $\alpha_d - (\gamma + 2\psi)D_d\sigma_d^2 = 0$. The binding constraint determines $\gamma + 2\psi = \sqrt{\mathrm{SR}_O^2 + \mathrm{SR}_U^2}/K$. The risk share:

$$w_O = \frac{D_O^2\sigma_O^2}{K^2} = \frac{\mathrm{SR}_O^2}{\mathrm{SR}_O^2 + \mathrm{SR}_U^2}.$$

Since $\mathrm{SR}_O < \mathrm{SR}_U$ for $t < t^*$ (Proposition 5), $w_O < 1/2$ for $t < t^*$, with equality at t^*. □

A.7 Proof of Proposition 7 (Welfare cost)

At $t = t^*$, $\mathrm{SR}_O = \mathrm{SR}_U \equiv \mathrm{SR}$. The asymmetric fund's alpha involves $\sqrt{\mathrm{SR}_O^2 + \mathrm{SR}_U^2} = \mathrm{SR}\sqrt{2}$. The symmetric fund's alpha involves $\sqrt{2}\mathrm{SR}_{\mathrm{sym}}(t^*)$.

Compute:

$$\frac{\mathrm{SR}_{\mathrm{sym}}(t^*)}{\mathrm{SR}} = \sqrt{\frac{\tau}{\tau_O}} \cdot \frac{1+1/\sqrt{k}}{2} = \frac{k^{1/4} + k^{-1/4}}{2},$$

where $\tau = \sqrt{\tau_O\tau_U}$. By the AM-GM inequality, $(k^{1/4} + k^{-1/4})/2 \geq 1$ with equality only when $k = 1$. Since $k > 1$, $\mathrm{SR}_{\mathrm{sym}}(t^*) > \mathrm{SR}$, so $\sqrt{2}\mathrm{SR}_{\mathrm{sym}}(t^*) > \sqrt{2}\mathrm{SR} = \sqrt{\mathrm{SR}_O^2 + \mathrm{SR}_U^2}$. The symmetric fund is still better at t^*, so the crossover $t_w > t^*$. □

A.8 Proof of Proposition 8 (Two-channel decomposition)

Part (a). With cost c_S, the short-leg alpha is $A_S(t) = m^O(t) + c_S = C/\pi_O(t) + c_S$. The long-leg alpha is $A_L(t) = |m^U(t)| = C/\pi_U(t)$. Then:

$$R_{TC}(t) = \frac{A_S}{A_L} = \frac{C/\pi_O + c_S}{C/\pi_U} = \frac{\pi_U}{\pi_O} + \frac{c_S \pi_U}{C} = R_I(t) + \frac{c_S \pi_U(t)}{C}.$$

Part (b). $S_I = R_I/(R_I + f)$ where $f(t) = c_S \pi_U(t)/C$ is increasing and linear. The derivative dS_I/dt has the sign of $R_I' f - R_I f'$. Compute $R_I'/R_I = \pi_0(\tau_U - \tau_O)/(\pi_O \pi_U)$ and $f'/f = \tau_U/\pi_U$. Then $R_I'/R_I < f'/f$ iff $\pi_0(1 - \tau_O/\tau_U) < \pi_O$, which holds for all $t \geq 0$. So $dS_I/dt < 0$. □

References

Alili, L., Patie, P., and Pedersen, J. L. (2005). Representations of the first hitting time density of an Ornstein-Uhlenbeck process. *Stochastic Models*, 21(4):967–980.

Ang, A., Hodrick, R. J., Xing, Y., and Zhang, X. (2006). The cross-section of volatility and expected returns. *Journal of Finance*, 61(1):259–299.

Bordalo, P., Gennaioli, N., La Porta, R., and Shleifer, A. (2019). Diagnostic expectations and stock returns. *Journal of Finance*, 74(6):2839–2874.

Bordalo, P., Gennaioli, N., La Porta, R., and Shleifer, A. (2024). Belief overreaction and stock market puzzles. *Journal of Political Economy*, 132(5).

Bordalo, P., Gennaioli, N., La Porta, R., and Shleifer, A. (2025). Finance without exotic risk. *Journal of Financial Economics*, forthcoming. NBER Working Paper 33004.

Brunnermeier, M. K. and Pedersen, L. H. (2009). Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6):2201–2238.

Chernov, M., Dahlquist, M., and Lochstoer, L. (2025). Unpriced risks: Rethinking cross-sectional asset pricing. NBER Working Paper 34009.

Daniel, K., Hirshleifer, D., and Sun, L. (2020). Short- and long-horizon behavioral factors. *Review of Financial Studies*, 33(4):1673–1736.

DeLong, J. B., Shleifer, A., Summers, L. H., and Waldmann, R. J. (1990). Noise trader risk in financial markets. *Journal of Political Economy*, 98(4):703–738.

Durbin, J. (1971). Boundary-crossing probabilities for the Brownian motion and Poisson processes and techniques for computing the power of the Kolmogorov-Smirnov test. *Journal of Applied Probability*, 8(3):431–453.

Gromb, D. and Vayanos, D. (2002). Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics*, 66(2-3):361–407.

Gu, S., Kelly, B., and Xiu, D. (2020). Empirical asset pricing via machine learning. *Review of Financial Studies*, 33(5):2223–2273.

Hong, H. and Kubik, J. D. (2003). Analyzing the analysts: Career concerns and biased earnings forecasts. *Journal of Finance*, 58(1):313–351.

Hou, K., Xue, C., and Zhang, L. (2020). Replicating anomalies. *Review of Financial Studies*, 33(5):2019–2133.

Jegadeesh, N. and Titman, S. (2001). Profitability of momentum strategies: An evaluation of alternative explanations. *Journal of Finance*, 56(2):699–720.

Kelly, B. T., Pruitt, S., and Su, Y. (2019). Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics*, 134(3):501–524.

Kozak, S., Nagel, S., and Santosh, S. (2020). Shrinking the cross section. *Journal of Financial Economics*, 135(2):271–292.

Linetsky, V. (2004). Computing hitting time densities for CIR and OU diffusions: Applications to mean-reverting models. *Journal of Computational Finance*, 7(4):1–22.

Lochstoer, L. A. and Tetlock, P. C. (2020). What drives anomaly returns? *Journal of Finance*, 75(3):1417–1455.

McLean, R. D. and Pontiff, J. (2016). Does academic research destroy stock return predictability? *Journal of Finance*, 71(1):5–32.

Ricciardi, L. M. and Sato, S. (1988). First-passage-time density and moments of the Ornstein-Uhlenbeck process. *Journal of Applied Probability*, 25(1):43–57.

Shleifer, A. and Vishny, R. W. (1997). The limits of arbitrage. *Journal of Finance*, 52(1):35–55.

Stambaugh, R. F. and Yuan, Y. (2017). Mispricing factors. *Review of Financial Studies*, 30(4):1270–1315.

Stambaugh, R. F., Yu, J., and Yuan, Y. (2012). The short of it: Investor sentiment and anomalies. *Journal of Financial Economics*, 104(2):288–302.

Stambaugh, R. F., Yu, J., and Yuan, Y. (2015). Arbitrage asymmetry and the idiosyncratic volatility puzzle. *Journal of Finance*, 70(5):1903–1948.

Van Binsbergen, J. H., Boons, M., Opp, C. C., and Tamoni, A. (2023). Dynamic asset (mis)pricing: Build-up vs. resolution anomalies. *Journal of Financial Economics*, 147(2):406–431.

Barber, B., Lehavy, R., McNichols, M., and Trueman, B. (2006). Buys, holds, and sells: The distribution of investment banks' stock ratings and the implications for the profitability of analysts' recommendations. *Journal of Accounting and Economics*, 41(1-2):87–117.

Da, Z., Nagel, S., and Xiu, D. (2024). The statistical limit of arbitrage. NBER Working Paper.

Fang, B., Hope, O.-K., Huang, Z., and Moldovan, R. (2020). The effects of MiFID II on sell-side analysts, buy-side analysts, and firms. *Review of Finance*, 24(4):855–902.

Garleanu, N., Panageas, S., and Zheng, G. (forthcoming). What to short. *American Economic Review*.

Guo, Y. and Mota, L. (2021). Should information be sold separately? Evidence from MiFID II. *Journal of Financial Economics*, 142(1):97–126.

Jegadeesh, N. and Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance*, 48(1):65–91.