RESEARCH NOTE RN-02 · SIGNAL CALIBRATION & QUORUM ECONOMETRICS

Non-Parametric Signal Calibration, Orthogonal Pillar Pooling & *One-Way Statutory Distress Vetoes.*

When scouts watch very different data feeds—insider stock purchases, quarterly cash-flow gaps, and supply-chain footnotes—you need a mathematical way to put them on the same scale, keep correlated scouts from double-counting the same news, and shrink unproven signals to zero.

AUTHOR: **STARLING QUANTITATIVE RESEARCH** PUBLISHED: **NOVEMBER 8, 2026** DESK: **STARLING QUANTITATIVE**

SECTION 01

Why Standard Bell-Curve Z-Scores & Simple Averages Break on Corporate Events

Once point-in-time disclosures are extracted without lookahead bias (Research Note RN-01), the next practical problem is combining them. An open-market insider purchase is measured in dollars, a cash-flow divergence is measured as a ratio of assets, and a customer-concentration shock is measured as a percentage of revenue. Before the safety layer can compare or combine those recommendations, they have to live on the same statistical scale.

Two common shortcuts cause multi-signal systems to misfire in live markets:

1. **One Outlier Breaks the Bell Curve:** Corporate events do not follow a neat bell curve; they are sparse and heavy-tailed. If you compute a standard Gaussian z-score $(x - \hat{\mu})/\hat{\sigma}$ over a trailing window, a single massive event blows out the sample standard deviation $\hat{\sigma}$ for months—making every normal, genuinely informative filing that follows look tiny by comparison.

2. **Counting the Same Filing Three Times:** When a company files a 10-Q showing strong operating cash flow and falling working-capital accruals, multiple accounting scouts will fire on the exact same afternoon. If you simply add their scores together, you treat three angles on the same balance sheet as if they were three independent coin flips—inflating conviction and over-sizing the bet.

SECTION 02

Domain-Split Non-Parametric Normalization: OOS Blom Quantiles & Zero-Anchored MAD

To achieve robustness against extreme outliers without in-sample distribution leakage, Starling Quantitative partitions scout feature domains into two non-parametric normalization families evaluated strictly over prior historical observations inside $[t_0,\ t_{\text{eval}} - 12\text{d}]$:

2.1 One-Sided Catalyst Magnitudes: Out-of-Sample Blom Half-Normal Quantiles

For strictly positive catalyst magnitudes $x_{k,i,t} > 0$ (such as clustered Form 4 officer purchases, 8-K contract awards, or regulatory approvals), we compute the tie-adjusted out-of-sample mid-rank $r_{k,i,t}$ against the N_{prior} strictly prior non-embargoed triggers and apply the Blom (1958) quantile transformation mapped to the folded half-normal distribution:

$$r_{k,i,t} = 1 + \sum_{m=1}^{N_{\text{prior}}} \mathbb{I}(x_m < x_{k,i,t}) + \frac{1}{2}\sum_{m=1}^{N_{\text{prior}}} \mathbb{I}(x_m = x_{k,i,t}), \qquad u_{k,i,t} = \frac{r_{k,i,t} - \frac{3}{8}}{N_{\text{prior}} + \frac{5}{4}} \in (0,1)$$

$$q^{+}(u_{k,i,t}) = \text{clip}\left(\Phi^{-1}\left(\frac{1+u_{k,i,t}}{2}\right),\ z^{+}_{\min},\ z^{+}_{\max}\right), \qquad 0 < z^{+}_{\min} < z^{+}_{\max}$$

Multiplying $q^{+}(u_{k,i,t})$ by the deterministic thesis sign $d_{k,i,t} \in \{-1, +1\}$ yields a bounded, outlier-invariant directional event score $z_{k,i,t} = d_{k,i,t} \cdot q^{+}(u_{k,i,t})$ (with collinear same-day multi-scout or multi-filing bursts inside the same economic domain dampened downstream in Step 03 via concave intra-pillar pooling).

2.2 Continuous Signed Series: Zero-Anchored MAD_0 Scaling

For continuous two-sided metrics $x_{k,i,t} \in \mathbb{R}$ where 0.0 is the natural economic origin (such as standardized unexpected earnings SUE, Sloan accrual ratios, or short-horizon return dislocations), we scale by the Gaussian-consistent **Zero-Anchored Median Absolute Deviation** ($\hat{s}_{0,k}$) floored by a fraction of the non-zero root-mean-square and a domain dispersion floor $s_{\text{floor},k} > 0$ to prevent division-by-zero explosions during low-dispersion regimes:

$$\hat{s}_{0,k} = 1.4826 \cdot \text{median}_{|x_m|>0}\big(|x_m|\big), \qquad s_{\text{eff},k} = \max\big(\hat{s}_{0,k},\ \ c_{\text{rms}}\ \text{RMS}_{|x_m|>0},\ \ s_{\text{floor},k}\big)\ ,$$

SECTION 03

One-Document-One-Witness Provenance, Concave Four-Pillar Pooling & Canary Glass Wall

Corporate catalysts decay at markedly different velocities depending on dissemination channel: an intraday SEC 8-K item or liquidity dislocation decays over days ($\tau_k = 5\text{d}$), whereas supply-chain footnote propagation or post-earnings drift persists over multi-week horizons ($\tau_k = 10\text{d}$). First, any same-sign signals sharing a normalized primary filing accession (source_artifact_id = SEC:<18-digit-accession>) collapse to the single highest-$|\tilde{z}|$ witness so multiple parsers reading one report never manufacture multi-scout confluence. Each normalized arrival impulse $z_{k,i,t_{\text{event}}}$ is then attenuated by discrete half-life decay $\tau_k \in \{5, 10\}\ \text{trading days}$ with a hard $3\tau_k$ horizon cutoff:

$$\tilde{z}_{k,i,t} = z_{k,i,t_{\text{event}}} \cdot 2^{-(t-t_{\text{event}})/\tau_k} \cdot \mathbb{I}(0 \le t - t_{\text{event}} \le 3\tau_k)$$

To prevent collinear scouts from double-counting a single economic shock, every scout and isolated sub-arm (arm_id) is mapped declaratively in our registry to one of **four orthogonal causal pillars** $p \in \{1, 2, 3, 4\}$:

- **Pillar 1 (P1) — Accounting & Insider Forensics:** Form 4 open-market officer/director purchase clusters and net-held option exercises (Code M), Form 8-K Accelerated Share Repurchase (ASR) dealer-covering contracts, Schedule 13D stakes, Item 1A novel numeric/covenant density, and XBRL Non-Cash Fortress vs. working-capital accruals.
- **Pillar 2 (P2) — Physical & Regulatory Exhaust:** Public federal regulatory clearances and recalls, administrative patent challenges, government procurement obligations, and catalyst-clock / product-safety telemetry.
- **Pillar 3 (P3) — Positioning & Quality Rebound:** Sector-relative solvency-screened $T + 3$ crash-stabilized liquidity rebounds on zero-distress balance sheets, index-rebalance flow shocks, and options volatility skew.
- **Pillar 4 (P4) — Network & Macro Plumbing:** Form 10-K ASC-280 shrunk customer-supplier transfer weights, single-stock follow-through-day quality compounder bids, and acute high-yield credit-spread impulse (HYG/IEF) monitors.

Within each pillar p, active positive shocks $\tilde{z}^{+}_{(1)} \ge \tilde{z}^{+}_{(2)} \ge \cdots > 0$ and negative shocks $\tilde{z}^{-}_{(1)} \le \tilde{z}^{-}_{(2)} \le \cdots < 0$ are sorted by magnitude and pooled via **symmetric two-arm concave summation** with diminishing marginal weight $\gamma_{\text{pool}} \in (0, 1)$ behind a **One-Way Canary/Shadow Glass Wall**—isolating probationary canary and shadow sub-arms (arm_id) so they never inflate GRADUATED core conviction:

$$S^{+}_{p,i,t} = \tilde{z}^{+}_{(1)} + \gamma_{\text{pool}} \sum_{j \geq 2} \tilde{z}^{+}_{(j)}, \qquad S^{-}_{p,i,t} = \tilde{z}^{-}_{(1)} + \gamma_{\text{pool}} \sum_{j \geq 2} \tilde{z}^{-}_{(j)}, \qquad S_{p,i,t} = S^{+}_{p,i,t} + S^{-}_{p,i,t}$$

Statutory restatement or delinquency filings (8-K Item 4.02, Form NT 10-K/10-Q), severe continuous distress shocks ($\tilde{z}_{k,i,t} \leq -z_{\text{veto}}$), or opposing-signal collisions ($S^{+}_{p,i,t} > 0 \wedge S^{-}_{p,i,t} < 0$) immediately clamp bullish conviction to 0.0. Across the four pillars, scores are combined via Empirical-Bayes-shrunk unit-simplex weights $\boldsymbol{\omega} \in \Delta^3$ ($\sum_{p=1}^{4} \omega_p = 1$, shrunk toward equal $1/4$ weights) and normalized by the uncentered arrival-event root-mean-square $\hat{\sigma}_S$ computed strictly over discrete catalyst arrivals $\mathcal{E}_{\text{active}}$:

$$Z_{i,t} = \text{clip}\left(\frac{\sum_{p=1}^{4} \omega_p \, S_{p,i,t}}{\max(\hat{\sigma}_S, \, \epsilon_S)}, \; -z_{\text{comp}}, \; +z_{\text{comp}} \right)$$

SECTION 04

Zero-Skill-Null Grinold Scaling & Post-Acceptance Exhaustion Haircuts

Composite conviction $Z_{i,t}$ is mapped into an expected H-day idiosyncratic residual return $\hat{\alpha}_{\text{gross},i,t}$ (in basis points) via the Grinold (1989) Fundamental Law of Active Management, scaling by the stock's clamped H-day residual volatility $\hat{\sigma}_{\text{resid},i,H}$ and the Empirical Bayes posterior Information Coefficient $\text{IC}_{\text{post},t}$:

$$\text{IC}_{\text{post},t} = \max\left(0.0, \quad \left(\frac{N_{\text{eff},t}}{N_{\text{eff},t} + \kappa_{\text{IC}}}\right) \widehat{\text{IC}}_{\text{OOS},t} + \left(\frac{\kappa_{\text{IC}}}{N_{\text{eff},t} + \kappa_{\text{IC}}}\right) \text{IC}_0 \right), \qquad \text{IC}_0 \equiv 0.0$$

Crucially, Starling fixes the prior strictly at the **zero-skill null** ($\text{IC}_0 = 0.0$) with linear Empirical Bayes shrinkage ($\kappa_{\text{IC}} > 0$) and one-sided zero truncation: small-sample signals are shrunk directly toward zero rather than buoyed by an optimistic heuristic prior, and any negative trailing rank correlation is truncated at 0.0.

Finally, because a filing accepted at τ_{accept} may already be partially priced by the market before batch order routing at t_{exec}, we deduct a **one-way Barra 2-Factor (Market + GICS Sector) post-acceptance price-exhaustion haircut** whenever the stock's idiosyncratic residual has already moved in the signal's direction between τ_{accept} and t_{exec}:

$$\hat{\alpha}_{\text{rem},i,t} = \hat{\alpha}_{\text{gross},i,t} \cdot \max\left(0,\ \ 1 - \frac{\max\left(0,\ d_{i,t} \cdot r_{i,\,\tau_{\text{accept}} \to t_{\text{exec}}}\right)}{\gamma_{\text{exh}}\ \hat{\sigma}_{\text{resid},i,1\text{d}}}\right)$$

An order is eligible for routing only if the remaining unpriced alpha clears expected round-trip execution friction by a strict safety multiple $\eta_{\text{cost}} > 1$: $\hat{\alpha}_{\text{rem},i,t} \geq \eta_{\text{cost}} \cdot \text{ExpectedRoundTripCost}_{i,t}$. Across our 5-fold Purged Walk-Forward evaluation on 3,607 Russell 3000 CIKs (2022–2026), replacing raw heuristic 1-to-10 scout confidence scores ($\text{Rank IC} = -0.0083$) with this calibrated four-pillar Grinold combiner lifted out-of-sample active-event 6-factor residual Rank IC to $+0.0901$ across the full admitted panel ($N = 794$ active events; 34 non-overlapping windows with mean $\text{Rank IC} = +0.1534, t = +2.26$), producing strictly monotonic net-of-cost return terciles ($+24.27$ bps $< +121.35$ bps $< +208.39$ bps) and $+144.57$ bps net residual alpha per event ($\text{PF} = 1.54, t_{\text{NW10}} = +2.05$) on the Graduated Core subset ($N = 218$ events across 54 non-overlapping cohorts).

SECTION 05

Deterministic One-Way Distress & Footnote Kill-Switch Vetoes

In early research iterations, a natural temptation is to place a generative language model committee at the end of a quantitative pipeline to review filing footnotes. In practice, generative advisory layers fail two institutional requirements: they introduce run-to-run sampling variance and contaminate historical walk-forward tests with pre-training lookahead bias. Starling Quantitative therefore excised all LLM components from the stack and enforces footnote and distress protection via **100% deterministic statutory and forensic kill-switches** under a strict **one-way veto invariant**:

$$w_{i,t}^{\text{final}} = \mathbb{I}_{\text{pass_veto}}(i,t) \cdot w_{i,t}^{\text{quant}}, \qquad \mathbb{I}_{\text{pass_veto}}(i,t) \in \{0,1\}$$

Whenever deterministic filing parsers detect a statutory accounting failure (Form 8-K Item 4.02 non-reliance restatements or Form NT 10-K/10-Q delinquency notices), acute 5-day high-yield credit-spread widening (HYG/IEF) on levered balance sheets, dedicated bearish forensic or hazard triggers (WARN Act plant closures, catalyst-clock binary hazards, federal product-safety recalls, or XBRL cash-burn and dilution traps), severe negative distress shocks ($\tilde{z}_{k,i,t} \leq -z_{\text{veto}}$), or simultaneous opposing-pillar collisions ($S^{+} > 0 \wedge S^{-} < 0$), $\mathbb{I}_{\text{pass_veto}}(i,t)$ evaluates to 0—clamping bullish allocation on our internal book to NO_TRADE (with zero single-stock inverse-ETF routing across all regimes) while preserving target_ticker (action_polarity = "SHORT_OR_LONG_VETO") in Feed Tier A so institutional long/short books with prime-broker borrow can express the short leg directly.

SECTION 06

Selected References & Technical FAQ

• Blom, G. (1958). *Statistical Estimates and Transformed Beta-Variables*. John Wiley & Sons.

• Grinold, R. C. (1989). The fundamental law of active management. *Journal of Portfolio Management*, 15(3), 30–37.

• Grinold, R. C., & Kahn, R. N. (1999). *Active Portfolio Management: A Quantitative Approach for Providing Superior Returns and Controlling Risk* (2nd ed.). McGraw-Hill.

• Huber, P. J. (1981). *Robust Statistics*. John Wiley & Sons.

Why do parametric sample mean and standard deviation z-scores fail on corporate filing and alternative data feeds?

Corporate event magnitudes exhibit extreme positive skewness and heavy tails. A single outlier inflates the parametric sample standard deviation $\hat{\sigma}$, compressing subsequent informative disclosures toward zero. Starling replaces Gaussian z-scores with Out-of-Sample Blom rank quantiles on one-sided catalysts and floored Median Absolute Deviation (MAD_0) scaling on two-sided continuous series.

How does concave 4-pillar pooling prevent correlated scouts from double-counting the same event?

When multiple scouts within the same economic domain trigger on a single corporate filing, linear summation overstates conviction. Starling first deduplicates same-sign signals sharing a primary filing accession (source_artifact_id), groups its scouts into four orthogonal causal pillars (P1–P4), and applies symmetric two-arm sorted concave pooling ($\tilde{z}_{(1)} + \gamma_{\text{pool}} \sum_{j \geq 2} \tilde{z}_{(j)}$) behind a One-Way Canary/Shadow Glass Wall before combining across pillars via Empirical-Bayes-shrunk unit-simplex weights normalized by arrival-event RMS.

Why does Starling shrink scout Information Coefficients toward a zero-skill prior ($\mathrm{IC}_0 = 0.0$)?

Shrinking noisy small-sample Information Coefficients toward a positive heuristic baseline manufactures positive expected alpha out of uninformative series. Starling shrinks every empirical residual IC strictly toward zero ($\mathrm{IC}_0 = 0.0$) via linear Empirical Bayes shrinkage with one-sided zero truncation so any signal with negative or uninformative out-of-sample rank correlation collapses to zero expected alpha.

NEXT RESEARCH NOTE · RN-03

© 2026 STARLING QUANTITATIVE · A BIWADA COMPANY

REGULATORY DISCLOSURE: Starling Quantitative is a proprietary quantitative research and internal capital allocation desk of Biwada LLC. Nothing in this monograph constitutes investment advice or an offer to solicit outside retail capital or manage third-party funds.