The Fisher Geometry of Signal Breadth: A Cramér–Rao Ceiling for the Information Ratio

Maksim Aleksandrovich Andreev, Joseph Howden
Quark Quantitative Research

Abstract

The fundamental law of active management (Grinold, 1989) states that the information ratio of an optimally combined signal set is IR = IC·√breadth, where breadth is the number of *independent* bets. With per-signal information coefficients largely fixed by the published anomaly literature, breadth is the binding constraint on attainable performance — yet "independent" and "breadth" are rarely made precise, and signal-combination machinery is rarely grounded in more than heuristic. We show that the central objects are exact estimation-theory quantities. For the cross-sectional return model $r = S\,a + \varepsilon$, $\varepsilon \sim N(0,\Sigma)$, with signals stacked columnwise in S: (i) the signal overlap matrix $M = S^T\Sigma^{-1}S$ *is* the Fisher information matrix of the loading estimate (an identity, not an analogy); (ii) the maximum attainable squared information ratio is the Cramér–Rao quadratic form $IR_{max}^2 = ic^T M^{-1}$ ic, of which Grinold's law is the orthonormal special case; (iii) effective breadth is the participation ratio and the exponential von Neumann entropy of the density matrix $\rho = M/tr(M)$, a genuine density matrix; and (iv) each signal's marginal contribution to the ceiling is a Schur complement, and the contributions of a forward sweep telescope *exactly* to the ceiling. These identities make the platform's density-matrix and Born-rule language a rigorous relabeling of Fisher information, not a metaphor, and yield a Cramér–Rao-grounded signal-admission criterion. Applied to a live 66-name market-neutral book spanning price-, options-, and earnings-derived signals, the framework places the achievable information ratio at a ceiling of approximately 0.094 per rebalance and shows — across three independently tested candidate signals — that these data domains are built out: further breadth requires a new, orthogonal data source, not another feature of the same domains. A distinct application of the same geometry — using it to set the capital risk budget across a set of already-confirmed return streams rather than to measure breadth — is shown, in a pre-registered out-of-sample test at small breadth (K = 3 near-orthogonal streams), to add nothing over equal-risk contribution, locating the construction's demonstrated value in diagnosis (the ceiling and a principled stopping rule) rather than in allocation. We give the construction (the Fisher-breadth module of the platform's engine library) with the four identities verified numerically, and the empirical ceiling on forward, point-in-time data per the validation posture of Andreev & Howden (2026e).

1. Introduction

Grinold's (1989) fundamental law, IR = IC·√breadth, isolates the two levers on risk-adjusted active return: the per-bet information coefficient (IC) and the number of independent bets (breadth). In a programme that draws its signals from the published cross-sectional anomaly literature — residual momentum (Blitz, Huij & Martens, 2011), betting-against-beta and low-volatility (Frazzini & Pedersen, 2014), the variance risk premium (Goyal & Saretto, 2009), the call–put implied-volatility spread (Cremers & Weinbaum, 2010), the volatility smirk (Xing, Zhang & Zhao, 2010), post-earnings drift (Bernard & Thomas, 1989) — the per-signal ICs are approximately fixed by what those papers report. The only lever the architect controls is *breadth*: how many genuinely independent signals the book combines.

This makes the precise measurement of breadth, and the rigour of the combination machinery, first-order questions. They are usually answered heuristically: breadth is taken as a raw signal count, "independence" as a pairwise-correlation rule of thumb, and the combination as an inverse-covariance weighting justified by appeal to Markowitz. We show that the correct objects are exact and already familiar from statistical estimation theory, and that — for a platform whose allocator speaks in density matrices and Born-rule amplitudes — they are the *same* objects, related by identity rather than by correspondence. The practical payoff is a Cramér–Rao information-ratio ceiling that says how much performance a given library can in principle deliver, and a signal-admission test that says when a library is built out and further effort should be redirected to a new data domain.

> **Contribution.** (i) We identify the signal overlap matrix used by the allocator with the Fisher information matrix of the signal-loading estimate — an exact identity. (ii) We show the attainable information ratio is the Cramér–Rao bound $ic^T M^{-1} ic$ and recover Grinold's law as its orthonormal special case. (iii) We identify effective breadth with the participation ratio and exponential von Neumann entropy of the density matrix ρ = M/tr(M), giving the platform's quantum language an exact, non-analogical reading. (iv) We give a Schur-complement decomposition of the ceiling that attributes attainable IR across signals and furnishes a scale-invariant admission criterion. (v) We report a live empirical breadth ceiling and a built-out verdict for the price/options/earnings domains. The construction and the numerical verification of all four identities are in the Fisher-breadth module of the platform's engine library.

2. The Signal-Combination Model

Consider a single rebalance over a cross-section of N assets. Let S be the N×K matrix that stacks K factor-neutralised signals columnwise (column s_k is signal k's score across the N assets, residualised against the risk factors so the book is dollar- and beta-neutral by construction), and let one-period excess returns follow the linear model

```
r = S a + ε ,    ε ~ N(0, Σ) ,
```

with a the K-vector of signal loadings and Σ the (estimation-cleaned) asset return covariance. In practice Σ is regularised by Marchenko–Pastur eigenvalue clipping (Laloux et al., 1999) or Ledoit–Wolf shrinkage (2004) before inversion; we treat Σ as given. The questions are: what linear combination of signals maximises the information ratio, what is the maximum attainable, and which signals contribute to it.

3. Four Identities

The objects below are exact consequences of the model in §2 — identities and standard estimation-theory results, not correspondences. Each is verified numerically in the construction's self-test of the Fisher-breadth engine module.

3.1 The signal overlap matrix is the Fisher information matrix

The Gaussian log-likelihood of the loadings is $L(a) = -\tfrac{1}{2}(r-Sa)^T\Sigma^{-1}(r-Sa) + \text{const}$. Its Fisher information — the curvature of the expected log-likelihood — is

$$I(a) = E[\, -\partial^2 L/\partial a \partial a^T \,] = S^T \Sigma^{-1} S =: M , \quad M_{kl} = s_k^T \Sigma^{-1} s_l .$$

So the overlap (Gram) matrix of the signals under the inverse-covariance metric — the object the Born-rule ensemble allocator already constructs to score signal redundancy — *is* the Fisher information of the signal-loading estimate (Fisher, 1922). Off-diagonal M_{kl} is the inner product of signals k and l in the Σ^{-1} metric; signals that point the same way through the covariance structure carry overlapping information and inflate the off-diagonal.

3.2 The information-ratio ceiling is the Cramér–Rao bound

The maximum squared information ratio over all linear combinations w = S c is

$$IR_{max}^2 = ic^T M^{-1} ic ,$$

where ic is the per-signal IC vector. Because M^{-1} is the Cramér–Rao lower bound on the covariance of any unbiased estimator of the loadings (Rao, 1945; Cramér, 1946), this quadratic form is the best attainable signal-to-noise the library can yield: no combination does better. The connection to Grinold is exact. When the signals are orthonormal in the Σ^{-1} metric (M = I) and share a common IC, $ic^T M^{-1} ic = K \cdot IC^2$, so $IR_{max} = IC \cdot \sqrt{K} = IC \cdot \sqrt{\text{breadth}}$ — the fundamental law, recovered as the orthonormal special case and correctly generalised to correlated signals by the inverse-Fisher form.

The IC-units normalisation. The IC of a signal is its correlation with returns — the edge of its *unit-*Σ^{-1}*-normalised* form. For $ic^T M^{-1} ic$ to be the squared information ratio, M must therefore be placed in correlation form M_c, with entries $M_{kl}/\sqrt{(M_{kk}M_{ll})}$ — the cosine similarity of the signals in the Σ^{-1} metric. Without this normalisation the raw quadratic form is expressed in mismatched covariance×signal units (of order 10^{-7} on daily returns) and is uninterpretable. Near-degenerate Fisher directions (two nearly collinear signals) are truncated before inversion, giving the robust ceiling and preventing the spurious blow-up in which the model "earns" infinite IR by trading the tiny difference of two redundant signals.

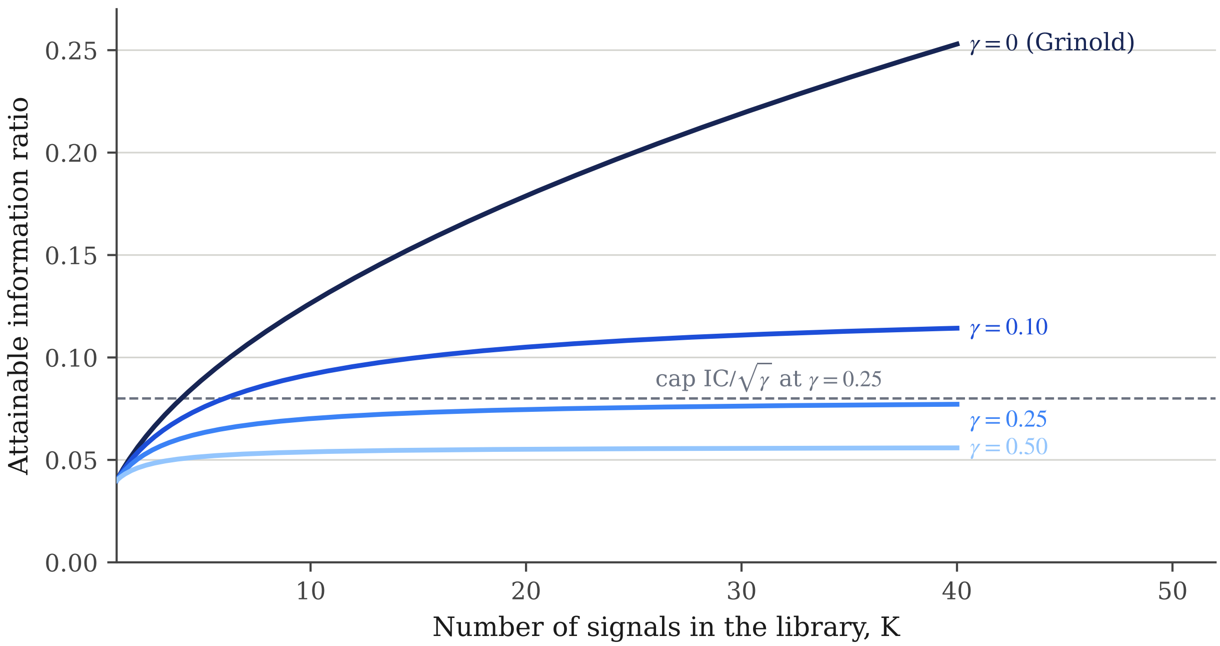

Figure 1. The Cramér–Rao ceiling generalises Grinold's law to correlated signals. Attainable information ratio $IR_{max} = IC \cdot \sqrt{(K / (1 + (K-1)\gamma))}$ against library size K, computed exactly from the equicorrelated form of the correlation-normalised Fisher matrix M_c (unit diagonal, common pairwise Fisher correlation γ) with a common per-signal IC of 0.04, mid-range of the published anomaly ICs of §1. The $\gamma = 0$ curve is Grinold's $IC \cdot \sqrt{K}$, growing without bound in the signal count; any $\gamma > 0$ caps the ceiling at $IC/\sqrt{\gamma}$ (dashed grey: 0.08 at $\gamma = 0.25$) no matter how many same-domain signals are added, and at $\gamma = 0.25$ roughly a dozen signals already come within 10% of the cap. Correlation, not signal count, is the binding constraint that the inverse-Fisher form makes exact.

3.3 Effective breadth is the participation ratio and the von Neumann entropy

Define the signal density matrix $\rho = M_c / tr(M_c)$. It is symmetric, positive semi-definite, and unit-trace — a genuine density matrix — whose eigenvalues p_i are the normalised Fisher spectrum. Two information-theoretic measures of the effective number of independent signals follow:

$$B_{PR} = (\mathrm{tr}\, M_c)^2 / \mathrm{tr}(M_c^2) = 1 / \Sigma_i\, p_i^2 \quad \text{(participation ratio)}$$

$$B_H = \exp(\, H(\rho)\,), \quad H(\rho) = -\Sigma_i\, p_i \ln p_i = -\mathrm{tr}(\rho \ln \rho) \quad \text{(entropy rank)}$$

Both equal K when the signals are Fisher-orthogonal ($\rho = I/K$) and collapse toward 1 as the signals become collinear. B_{PR} is the inverse-Simpson (Herfindahl) concentration of the Fisher spectrum; B_H is the Roy–Vetterli (2007) entropy effective rank, and $H(\rho)$ is exactly the von Neumann entropy of the density matrix. This is the precise, non-analogical sense in which the platform's quantum language applies: ρ is literally a density matrix and $H(\rho)$ is literally its von Neumann entropy; the “effective number of independent strategies” the Born-rule allocator reports is the exponential of that entropy.

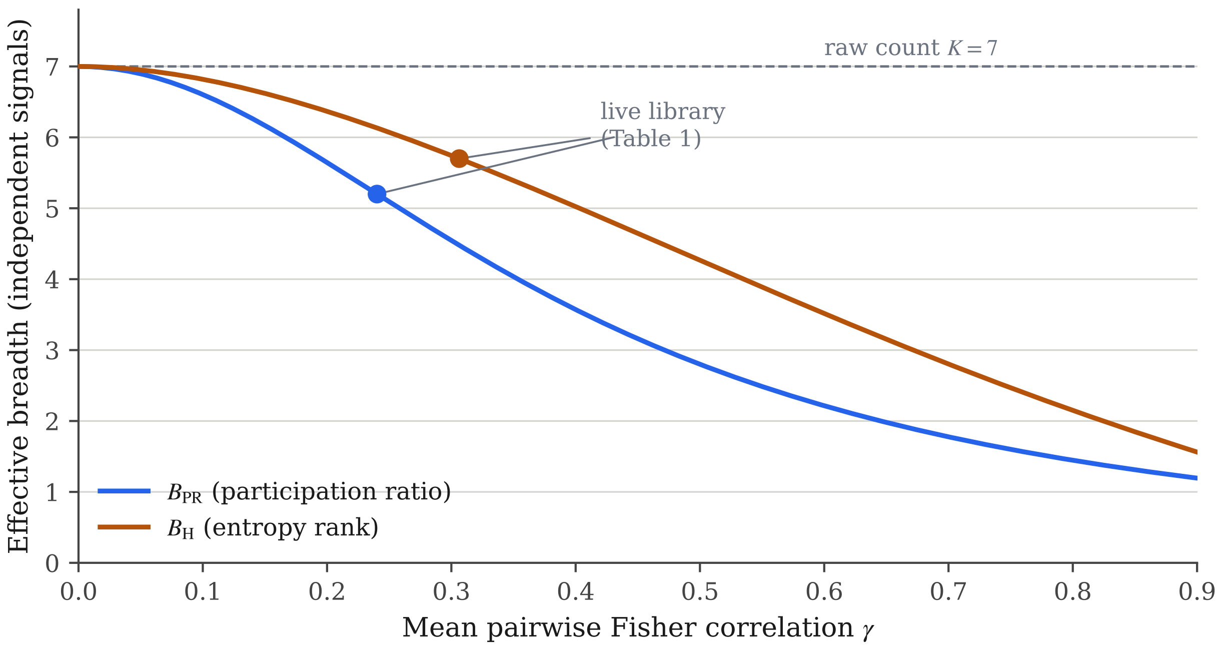

Figure 2. Effective breadth collapses under signal overlap. Participation breadth $B_{PR} = K / (1 + (K-1)\gamma^2)$ (blue) and entropy-rank breadth $B_H = \exp H$ (amber) for a $K = 7$ library whose signals share a common pairwise Fisher correlation γ, computed exactly from the two eigenvalues $1 + (K-1)\gamma$ and $1 - \gamma$ of the equicorrelated M_c. Both equal the raw count 7 (dashed grey) only at $\gamma = 0$ and collapse toward 1 as the library collinearises, the entropy rank being the more forgiving of the two counts at moderate overlap. The markers place the live seven-signal library of Table 1 on the model curves: its measured $B_{PR} \approx 5.2$ and $B_H \approx 5.7$ are reproduced at $\gamma \approx 0.24$ and $\gamma \approx 0.31$ respectively. The two implied correlations differ because the live Fisher spectrum is not exactly equicorrelated; together they bracket the library's mean pairwise overlap at roughly a quarter to three-tenths.

3.4 Each signal's marginal value is a Schur complement

The increment to IR_{max}^2 from adjoining signal k to an already-included set A is the Schur complement of the bordered Fisher matrix:

$$\Delta_k = (\, ic_k - m_{Ak}^T M_{AA}^{-1}\, ic_A\,)^2 \;/\; (\, m_{kk} - m_{Ak}^T M_{AA}^{-1}\, m_{Ak}\,)$$
$$= (\text{Fisher-orthogonalised IC of } k)^2 / (\text{Fisher-orthogonal residual norm of } k)\,.$$

The numerator is signal k's IC after projecting out its overlap with the included signals; the denominator is the norm of its Σ^{-1}-orthogonal residual. A signal that is redundant (small residual norm) or carries no orthogonal edge (small orthogonalised IC) contributes

≈ 0. The increments of a complete forward sweep telescope *exactly* to the full $ic^T M^{-1} ic$ — an exact additive attribution of the attainable information ratio across signals (verified to machine precision). This is the principled, scale-invariant admission criterion that pairwise-correlation and leave-one-out heuristics only approximate.

Admission rule. A candidate signal is admitted only if its Schur increment exceeds a scale-invariant floor — a fixed fraction of the library's total attainable IR^2. Individual statistical significance is deliberately *not* the gate: a book's value is that the combination of many weak-but-real signals clears, and each signal's own significance is settled by the forward Harvey test (Harvey, Liu & Zhu, 2016), not by in-sample selection.

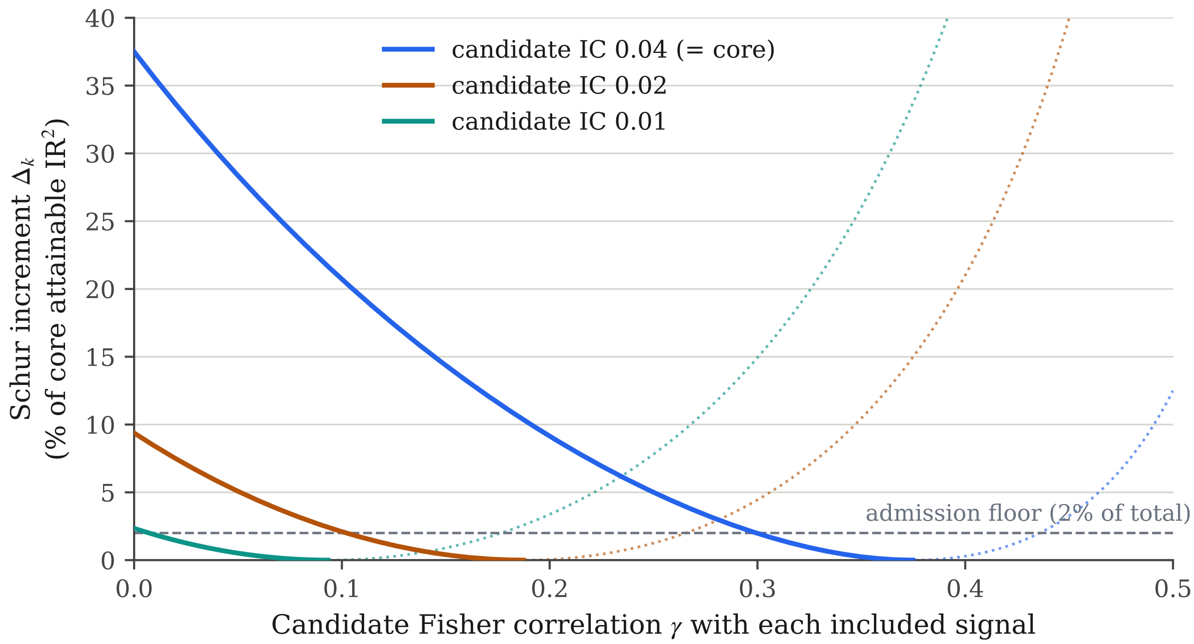

Figure 3. The Schur admission gate of §3.4 in operation. Marginal increment Δ_k from adjoining one candidate signal to an included six-signal core (equicorrelated at $\gamma = 0.25$ with common IC 0.04, approximately the geometry of the §6 library), expressed as a share of the core's total attainable IR^2 and computed exactly from the Schur-complement identity, against the candidate's Fisher correlation γ with each included signal, for candidate ICs of 0.04 (equal to the incumbents), 0.02 and 0.01. The dashed grey line is the scale-invariant admission floor (2% of the total attainable IR^2). A full-edge candidate falls below the floor once its overlap with the core exceeds $\gamma \approx 0.30$; a half-edge candidate at $\gamma \approx 0.10$; a quarter-edge candidate sits essentially at the floor even when fully orthogonal. Each increment vanishes where the candidate's Fisher-orthogonalised IC crosses zero ($\gamma = 0.375$ for the full-edge candidate); the thin dotted branches beyond are the residual-contrarian region, in which the identity credits the candidate only as a short hedge against the core — the gate does not admit signals there (the §6 lead-lag candidate, with negative orthogonalised IC, was dropped, not shorted), and the near-degenerate truncation of §3.2 prevents that branch from amplifying into the collinear blow-up. This is the mechanism behind the §6 marginal verdicts: redundancy, not absence of edge, is what excludes a candidate.

4. The Density-Matrix Relabeling Is Exact

The platform's allocator is described elsewhere (Andreev & Howden, 2026a) in the language of quantum measurement: signals carry Born-rule amplitudes, the ensemble is weighted by $|\Phi|^2$, and diversity is an “effective number of strategies.” The results of §3 show this language is not a suggestive analogy but an exact relabeling. The matrix $M = S^T \Sigma^{-1} S$ is simultaneously (a) the Gram matrix that defines signal overlap, (b) the Fisher information of the estimation problem, and (c), once unit-normalised and trace-normalised, a density matrix ρ. Its spectrum is the menu of independent information directions; its von Neumann entropy is the diversity; its inverse quadratic form against the IC vector is the attainable information ratio. The Born-rule-style weighting $w \propto M^{-1}ic$ that the allocator applies is the Fisher-efficient (generalised-least-squares) estimator of the loadings — the minimum-variance unbiased combination. Every term in the quantum description corresponds, by a provable identity, to a term in the estimation-theoretic description. This is the standard the programme holds itself to: a correspondence is admitted only when it is a theorem, and discarded as decoration otherwise.

5. Operational Use

The construction is used in three ways. First, IR_{max} is the honest target: we report not a backtested Sharpe but "the library can attain at most IR_{max}; the live book is realising a fraction C of it," with the forward shadow curve measuring C out of sample. Second, the Schur increments drive Fisher-efficient selection: signals below the scale-invariant floor are flagged as adding breadth but no incremental information ratio. Third, and most consequentially, the increments form a *diminishing-returns detector*: when successive candidate signals all return sub-floor increments, the library is built out for its data domains, and the correct response is to stop adding signals and acquire a new, orthogonal data source. This is the value the framework adds beyond a raw breadth count — it states, rigorously, when to stop.

6. Empirical Breadth Ceiling

We assemble a live signal library for a 66-name, dollar- and beta-neutral US large-cap book, drawing from three data domains: price (residual momentum, short-term reversal, low-volatility), options-implied (Cremers–Weinbaum call–put IV spread, Goyal–Saretto variance-risk premium, Xing–Zhang–Zhao smirk), and earnings (post-earnings drift). The covariance is Marchenko–Pastur-cleaned; signals are factor-neutralised before forming M.

QUANTITY	VALUE
Signals in library	7
Participation breadth B_{PR}	≈ 5.2
Entropy-rank breadth B_H	≈ 5.7
Cramér–Rao IR ceiling (per rebalance)	≈ 0.094
Fisher-efficient core	≈ 6 signals

The decisive empirical finding concerns marginal value. Three candidate signals were each tested through the Schur-increment gate and each found marginal: an intra-industry lead-lag signal (Hou, 2007) adds breadth but negative orthogonalised IC (its catch-up content is already carried by short-term reversal); and the second and third options-implied signals (the call–put spread and the smirk) each raise the ceiling only slightly, because the variance-risk-premium signal already captures the bulk of the options-implied dimension. The pattern is consistent and was confirmed three times independently:

> **Built-out verdict.** The price, options, and earnings domains are built out to a breadth ceiling of approximately 0.094 per-rebalance information ratio. Within these domains each additional signal is marginal. Materially raising the ceiling requires a *genuinely new, orthogonal data domain* — not another feature of the domains already mined. The candidate frontier the framework points to is market microstructure / dealer positioning (the self-consistent dealer-flow field of Andreev & Howden, 2026f) — a data domain whose information is, by construction, uncorrelated with price-, options-, and earnings-derived signals, and which is currently accruing on forward, point-in-time data with no breadth verdict yet rendered. Whether it in fact raises the ceiling is an open, forward question, not a result claimed here.

> **The anchor book's gross alpha has since been measured, and it corroborates the built-out verdict from the other side.** The realised fraction referred to in the caveats below is no longer entirely open. Assembled and run, the same 66-name book returns a gross alpha of +0.041% per year with a Newey–West standard error at five lags giving $t = +0.013$, which is statistically zero. A layer-by-layer attribution of that result across the construction rungs — neutralisation, covariance cleaning, combination weights, sizing — found no rung responsible and returned the construction stack exonerated, locating the binding constraint in the signal library itself. This is the paper's thesis measured from the opposite direction: if the domains are built out at a ceiling of roughly 0.094 and the assembled book earns nothing distinguishable from zero, the deficit is in what the signals know rather than in how they are combined, and adding another feature of an already-mined domain will not change it. We report the figure because omitting a measured zero from a paper about attainable ceilings would leave the strongest available corroboration unstated.

> **Caveats.** The ceiling is an attainable bound under the model of §2 with ICs taken from the literature; the realised fraction measured above is one book over one window, not a general attainment rate. ICs estimated in-sample would inflate the ceiling, which is why the admission gate uses relative contribution rather than in-sample significance and the realised performance is measured forward. The covariance regularisation choice (Marchenko–Pastur vs shrinkage) perturbs M and is reported as a sensitivity.

7. Boundary of Validity: Risk Budgeting at Small Breadth

The identities of §3 validate the construction as a *diagnostic*: they make the breadth count, the attainable ceiling, and the stopping rule exact. A separate and stronger question is whether the same Fisher geometry is useful *allocatively* — whether the Schur increments, which attribute the ceiling across streams, should also *set the capital risk budget* across them. We pre-registered and tested this directly, and report the result here so the diagnostic claim is not silently over-extended into an allocation claim it does not support.

The test allocates the risk budget across three already-confirmed, near-orthogonal return streams — a trend-following harvester, an index variance-risk-premium income harvester, and an event-vega harvester — over 119 walk-forward out-of-sample months (2016–2026). The information-geometric budget sets each stream's risk share from its Schur increment Δ_k (its marginal contribution to the squared ceiling, using each stream's own trailing realised information as its IC), and is compared against equal-risk contribution (ERC; Maillard, Roncalli & Teiletche, 2010) and against a size-matched control of 2,000 random Dirichlet risk budgets. Every book holds the same three streams under the same covariance and the same inverse-volatility scaling; only the budget vector differs, which isolates the budgeting rule.

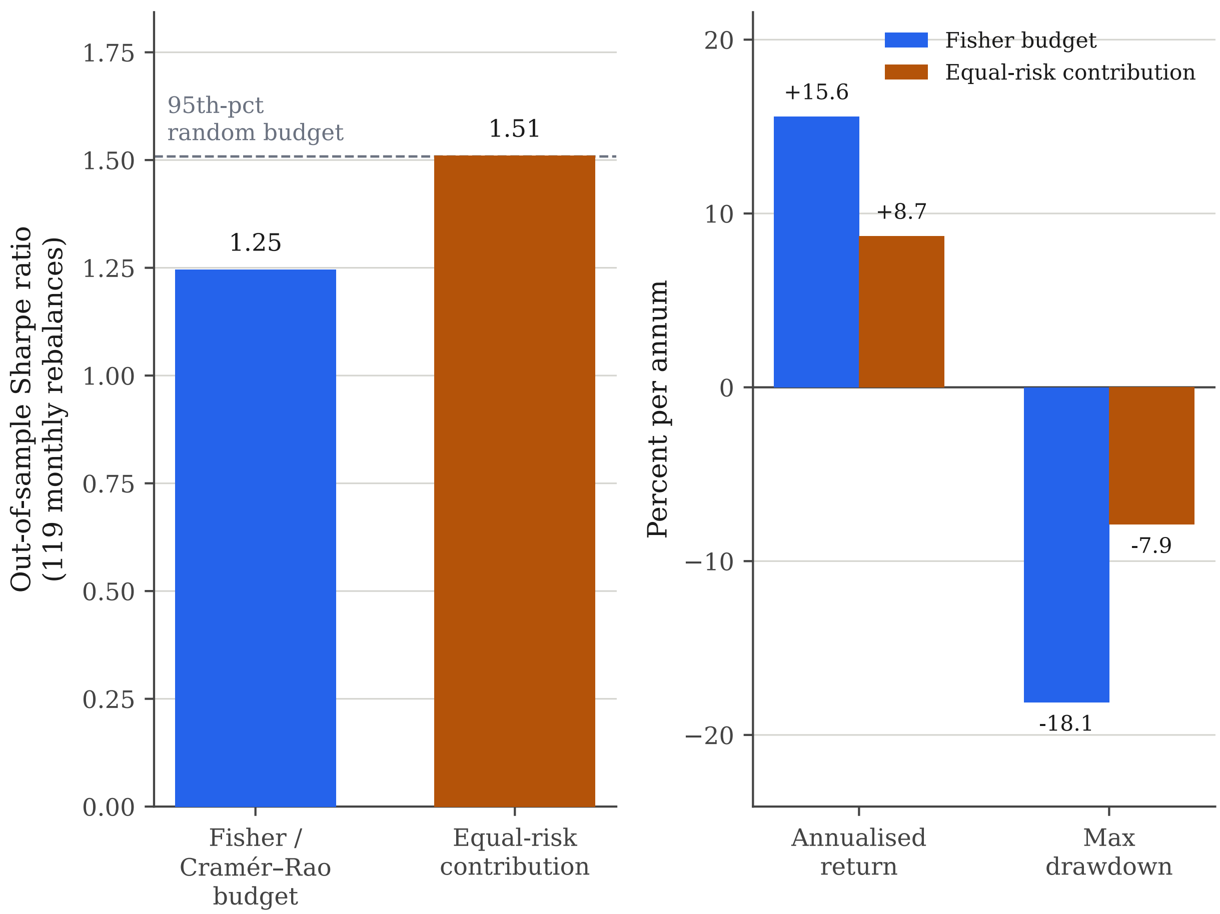

Figure 4. Information-geometric risk budgeting is inert at small breadth. Left: out-of-sample Sharpe ratio over 119 monthly rebalances (2016–2026) for a book of three confirmed, near-orthogonal return streams, allocated by the Fisher / Cramér–Rao Schur-increment risk budget (blue) versus equal-risk contribution (amber); the dashed line marks the 95th percentile of 2,000 random Dirichlet risk budgets. The Fisher budget (Sharpe 1.25) falls below both equal-risk contribution (1.51) and the random-budget 95th percentile — it ranks at only the 31st percentile of the random control, and the scored difference is −0.26 Sharpe. Right: the Fisher budget earns a higher annualised return (15.6% vs 8.7%) only by concentrating into a roughly 2.3× deeper maximum drawdown (−18.1% vs −7.9%), which is why its Sharpe is lower. With three near-orthogonal streams of similar information ratio, equal-risk contribution already sits on the Cramér–Rao ceiling and there is no redundancy for the Schur increments to exploit. Source: the platform's pre-registered risk-budget control study (frozen 2026-06-24).

The verdict is a settled sub-bar null. Three of four pre-registered bars fail: the paired Newey–West t on the monthly return difference is 2.0 (below the required 3.0); the Fisher book's Sharpe of 1.25 sits at only the 31st percentile of the random-budget control rather than beating its 95th; and, contrary to the control-system hypothesis, the Fisher book *concentrates further*, not less, in the low-breadth months. The mechanism is the pre-recorded one and is decisive: when the streams are near-orthogonal with similar information ratios, the maximum-information-ratio combination *is* equal risk, so ERC already sits on the Cramér–Rao ceiling — its budget is uniform every month, at full breadth. The Schur increments then merely amplify differences in trailing realised Sharpe, tilting toward whichever stream ran hottest and buying a deeper drawdown for no incremental information.

> **Diagnostic, not allocative — at this breadth.** The construction is validated as a measurement instrument (the ceiling, the two breadth counts, the exact Schur attribution, the stopping rule), not as a risk-allocation rule. The control-system reading is shown *inert* at K = 3 near-orthogonal confirmed streams, where there is no redundancy for the Cramér–Rao ceiling to exploit — it is not refuted in general. Whether information-geometric budgeting adds value at larger libraries with genuine redundancy, or for streams that collinearise in crises, is the scoped open case; the same telescoping Schur identity that fails to help here is exactly the object that would locate and de-duplicate that redundancy when it exists.

8. Discussion

Breadth is the binding constraint on attainable active return, and we have shown it is an exact estimation-theoretic quantity: the participation ratio and von Neumann entropy of a Fisher information matrix that is simultaneously the signal overlap matrix and a density matrix. The attainable information ratio is the Cramér–Rao bound, with Grinold's law its orthonormal limit. This furnishes both a ceiling — how good a given library can be — and a stopping rule — when a library is built out and effort should move to a new data domain. Applied live, the framework places a 66-name multi-domain book at a per-rebalance IR ceiling near 0.094 and returns a built-out verdict for the price/options/earnings domains, redirecting breadth-growth effort to microstructure. More broadly, the result is an instance of the programme's standard: the density-matrix and Born-rule language of the allocator is not borrowed imagery but the same mathematics as Fisher information and the Cramér–Rao bound, related by identity and verified numerically. Empirical realisation of the ceiling is reported separately and only on forward data, per Andreev & Howden (2026e).

References

Bernard, V. L. & Thomas, J. K. (1989). Post-earnings-announcement drift: delayed price response or risk premium? *Journal of Accounting Research*, 27, 1–36.

Blitz, D., Huij, J. & Martens, M. (2011). Residual momentum. *Journal of Empirical Finance*, 18(3), 506–521.

Cramér, H. (1946). *Mathematical Methods of Statistics*. Princeton University Press.

Cremers, M. & Weinbaum, D. (2010). Deviations from put-call parity and stock return predictability. *Journal of Financial and Quantitative Analysis*, 45(2), 335–367.

Fisher, R. A. (1922). On the mathematical foundations of theoretical statistics. *Philosophical Transactions of the Royal Society A*, 222, 309–368.

Frazzini, A. & Pedersen, L. H. (2014). Betting against beta. *Journal of Financial Economics*, 111(1), 1–25.

Goyal, A. & Saretto, A. (2009). Cross-section of option returns and volatility. *Journal of Financial Economics*, 94(2), 310–326.

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

Harvey, C. R., Liu, Y. & Zhu, H. (2016). …and the cross-section of expected returns. *Review of Financial Studies*, 29(1), 5–68.

Hou, K. (2007). Industry information diffusion and the lead-lag effect in stock returns. *Review of Financial Studies*, 20(4), 1113–1138.

Laloux, L., Cizeau, P., Bouchaud, J.-P. & Potters, M. (1999). Noise dressing of financial correlation matrices. *Physical Review Letters*, 83(7), 1467–1470.

Ledoit, O. & Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2), 365–411.

Maillard, S., Roncalli, T. & Teiletche, J. (2010). The properties of equally weighted risk contribution portfolios. *Journal of Portfolio Management*, 36(4), 60–70.

Markowitz, H. (1952). Portfolio selection. *Journal of Finance*, 7(1), 77–91.

Rao, C. R. (1945). Information and the accuracy attainable in the estimation of statistical parameters. *Bulletin of the Calcutta Mathematical Society*, 37, 81–91.

Roy, O. & Vetterli, M. (2007). The effective rank: a measure of effective dimensionality. *European Signal Processing Conference (EUSIPCO)*, 606–610.

Xing, Y., Zhang, X. & Zhao, X. (2010). What does the individual option volatility smirk tell us about future equity returns? *Journal of Financial and Quantitative Analysis*, 45(3), 641–662.

Andreev, M. A. & Howden, J. (2026a). Stochastic mechanical methods for quantitative portfolio management. *Quark Quantitative Research Working Paper*.

Andreev, M. A. & Howden, J. (2026e). Ontological limits of historical validation for wavefunction-calibrated financial systems. *Quark Quantitative Research Working Paper*.

Andreev, M. A. & Howden, J. (2026f). A self-consistent dealer-flow field: nonlinear-Schrödinger modelling of spot–options feedback. *Quark Quantitative Research Working Paper*.