

**The multiplicative index of riskiness in the cross-section of U.S. equity returns:
Theory and evidence for 1928-2025**

Tom Miller
Professor of Finance
Michael J. Coles College of Business
Kennesaw State University
(470) 578-6494
tmiller@kennesaw.edu
Corresponding author

The multiplicative index of riskiness in the cross-section of U.S. equity returns: Theory and evidence for 1928-2025

Abstract

Finance is multiplicative. An investor with wealth w who holds an asset of gross return φ ends with $w\varphi$, yet the objective riskiness of assets is almost always measured with the additive, dollar-denominated Aumann-Serrano (2008) index, calibrated through constant absolute risk aversion. Its constant-relative-risk-aversion (CRRA) counterpart, due to Li (2014a) and Schreiber (2014), is the theoretically natural instrument for asset returns but has a thin empirical record. We develop this multiplicative index into a usable empirical instrument and apply it to the size and value dimensions of the U.S. equity cross-section for 1928-2025.

The theory is threefold: a two-moment closed form under log-normal returns; an existence characterization showing that the multiplicative and additive indices read opposite tails of the return distribution, so that heavy right tails leave the multiplicative index finite and closed-form where the additive index of a short position diverges; and horizon invariance for i.i.d. returns, obtained through the logarithm. Empirically, across 27 assets and five periods the two indices are nearly ordinally equivalent, yet their reversals localize among the value deciles. The central finding is that the critical CRRA coefficient is non-monotone across the book-to-market sort, humped in the interior, monotone across size, so deep value is unattractive to a moderately risk-averse agent despite the book's highest mean log return, through variance drag alone. A properly powered monthly test shows this ordering is not priced: the index is an instrument for measuring riskiness and screening portfolios rather than for pricing the cross-section.

Keywords: riskiness; relative risk aversion; multiplicative gambles; CRRA; size and value; cross-section of returns; performance measurement

JEL classification: G11, G12, D81, C58

1. Introduction

Whether an investor holds an asset depends on two separate things: the attributes of the asset and the attributes of the investor. Arrow (1965) and Pratt (1964) made the second precise with the coefficient of risk aversion, a subjective quantity that lives in the investor's utility function. They left the first unaddressed. Aumann and Serrano (2008) supplied it, defining an *objective index of riskiness*: a number attached to a gamble alone, independent of who evaluates it. Their index $R(g)$ is the reciprocal of the constant absolute risk aversion (CARA) at which an agent is indifferent to the gamble, and it is characterized by two axioms, a duality condition asserting that uniformly less risk-averse agents accept riskier gambles, and positive homogeneity.

The Aumann and Serrano index is additive: it is defined on payoffs g added to wealth, and it is denominated in dollars. Finance, however, is multiplicative. An investor with wealth w who buys an asset with gross return φ ends with $w\varphi$, not $w + g$, and the natural calibrating class is not CARA but constant relative risk aversion (CRRA). Li (2014a) and, four months later and independently, Schreiber (2014) built the counterpart: an index defined on multiplicative gambles in which relative risk aversion plays the role absolute risk aversion plays in the additive

theory. We call it the *multiplicative* (or *relative*) index and write it R_m , with critical CRRA coefficient $\alpha_m = 1/R_m$.

This is the theoretically natural object for asset returns, and it is almost entirely unused. The applied literature built on Aumann and Serrano is additive: the index is deployed chiefly as a performance measure that generalizes the Sharpe ratio to non-normal returns (Homm and Pigorsch, 2012; Kadan and Liu, 2014). We are aware of a single empirical application of the multiplicative index, a time-series test of time diversification on three equity indices by Lu, Yang and Wong (2018), and of no cross-sectional application at all. The theoretically preferred instrument is the one the empirical literature has left on the shelf. This paper takes it down and puts it to work on the canonical laboratory of empirical asset pricing: the size and value cross-section of U.S. equities.

Two obstacles have kept the multiplicative index off the shelf and the paper removes both. The first is computational. The index is defined implicitly by a root of an integral equation and for several relevant return distributions the root does not exist. We show (Section 3) that this problem essentially evaporates for asset returns. Under the Black-Scholes-consistent benchmark that log returns are normal, the index has an elementary closed form in two moments,

$$\alpha_m(\varphi) = 1 + 2m/s^2, \quad R_m(\varphi) = s^2/(s^2 + 2m), \quad (1)$$

and it always exists. More importantly, we prove a characterization that explains *why*: the multiplicative and additive indices read opposite tails of the return distribution. The multiplicative index sees only the left tail of the log return; the additive index of a short position sees only the right tail of the gross return. Heavy right tails, the canonical case for gross returns, and exactly what makes the additive short index diverge and forces the extension of Schulze (2014), leave the multiplicative index untouched. There is a converse relationship: on distributions with a heavy left tail and a light right tail, the additive index exists while the multiplicative one does not. This turns the paper's organizing intuition, that the multiplicative index, usually treated as a derived curiosity, is the better-behaved object on financial returns, from rhetoric into a theorem with a clean domain statement.

The second obstacle is that no one has shown the multiplicative index tells us anything an additive measure would not. Here the paper's central empirical finding does the work. Computing both indices for 27 assets over five non-overlapping periods from 1928 to 2025, we find that the two indices are very nearly ordinally equivalent Spearman correlations between 0.995 and 0.999, with only 12 of 1,755 asset pairs discordant; so at the level of a full-panel ranking the choice of index rarely matters. But the reversals are not random: they occur among the value deciles, and they are the visible edge of a structural fact. The critical CRRA coefficient α_m is *non-monotone in the book-to-market decile*. It rises from growth to the middle of the book and falls again into deep value, tracing an interior hump; a quadratic fit places the vertex at decile 4.88, an interior maximum, whereas the same fit on the size book places the vertex at 10.01, on the boundary. The two extremes of the value book carry a critical CRRA coefficient 32.6% below the book's interior ($t = -3.33$, $p = 0.006$).

The economics of that hump is the paper's payoff and it is visible only through the CRRA lens. Because α_m is monotone in the single quantity $2m/s^2$, its units are a relative risk aversion coefficient, and every value reads directly as a statement about who will hold the asset: a CRRA agent holds it if and only if $\gamma < \alpha_m$. Growth stocks sit low on that scale for the ordinary reason, a

low mean log return. Deep-value stocks sit just as low *despite having the highest mean log return in the book*, because their variance drag $s^2/2$ is large enough to pull the risk-adjusted quantity $2m/s^2$ back down. A dollar-denominated additive measure, or a raw dispersion measure, records deep value as high-mean and volatile; the CRRA-calibrated index records the two effects on the one scale that decides acceptance and reveals that both ends of the value book are unattractive to a moderately risk-averse agent, for opposite reasons. This is a statement about the value premium, about the shape of its riskiness across the book, that a dollar-denominated measure cannot make in these terms.

A structural feature of the cross-section invites a pricing question: if the critical CRRA coefficient humps across the value book, does that shape show up in expected returns? We answer it directly, and the answer is no. Moving to monthly data, twenty size and value deciles, roughly a thousand cross-sections, a Fama-MacBeth test finds that trailing relative riskiness does not price the equity cross-section (a flat slope, $t \approx 0$), that a riskiest-minus-safest decile portfolio earns essentially nothing, and that an apparent annual link between the hump and subsequent value-decile returns dissolves once the test is properly powered. The hump is a descriptive property of the return distribution, not a priced factor; average returns across these deciles are governed by the size and value premia, which the multiplicative index does not encode. We report this null as a finding rather than bury it: it fixes the index's role as an instrument of measurement and portfolio screening rather than cross-sectional pricing, and it forecloses an over-reading of the hump that the annual panel alone would have invited.

The paper's contributions are therefore four. (i) *Theory*: a two-moment closed form, an existence characterization phrased as a tail dichotomy with a converse, and a horizon-invariance-and-products result obtained through the logarithm (Sections 2-3). (ii) *Method*: the multiplicative index turned into a reproducible empirical instrument for the size and value cross-section (Sections 4-5). (iii) *Descriptive evidence*: a full cross-sectional application in which the non-monotonicity of the critical CRRA coefficient across the value book is the signature result, supported by rank-agreement, CRRA-band, dependence, horizon, and non-parametric analyses (Sections 6-8). (iv) *A pricing null*: a properly powered monthly test, Fama-MacBeth, portfolio sorts, and a GMM stochastic-discount-factor test, showing that the riskiness ordering, and the hump within it, are neither predictive of nor priced in the size and value cross-section (Section 9). Section 10 concludes. Throughout, we report the critical CRRA coefficient $\alpha_m = 1 + \alpha(\log \varphi)$ of Li rather than the relative riskiness of Schreiber; the two are monotone transforms and carry the same ordinal content, but α_m reads directly as a risk aversion threshold.

2. The two indices and their exact relationship

This section fixes notation and states the one identity that drives everything downstream. A *utility function* is strictly increasing, strictly concave and twice continuously differentiable on $\mathbb{R}$. An *additive gamble* g is a random variable with $Eg > 0$ and $P(g < 0) > 0$. Following Aumann and Serrano, the riskiness $R(g)$ is the reciprocal of the unique *attractiveness* $\alpha(g) > 0$ solving

$$E[e^{-\alpha(g)g}] = 1. \quad (2)$$

The map $\beta \mapsto E[e^{-\beta g}] - 1$ is convex, vanishes at 0 with derivative $-Eg < 0$, and tends to $+\infty$ because g takes negative values; the positive root therefore exists and is unique (Aumann and Serrano, 2008; Schulze, 2014, for the domain question). Theorem B of Aumann and Serrano fixes the interpretation: $R(g)$ is the reciprocal of the CARA coefficient at which an agent is

indifferent to g , and CARA is the unique class whose members accept or reject a gamble independently of wealth, the wealth-free medium an objective index requires.

A *multiplicative gamble* φ is a strictly positive random variable with $E \log \varphi > 0$ and $P(\varphi < 1) > 0$, exactly the conditions that make $g = \log \varphi$ an additive gamble. An agent with wealth w who takes φ ends with $w\varphi$ and accepts if $Eu(w\varphi) > u(w)$. Recall that a CRRA utility with parameter $\gamma > 0$ has relative risk aversion $r(w) = -wu''(w)/u'(w) \equiv \gamma$. Li (2014a) proves the pivotal acceptance result: a CRRA agent with $\gamma > 1$ accepts φ at every wealth if and only if $\gamma < 1 + \alpha(\log \varphi)$, while every agent with $\gamma \leq 1$ accepts every admissible φ . The critical CRRA parameter defines the relative attractiveness and relative index

$$\alpha_m(\varphi) = 1 + \alpha(\log \varphi), \quad R_m(\varphi) = 1/\alpha_m(\varphi), \quad (3)$$

equivalently, $\alpha_m(\varphi)$ is the unique number exceeding 1 with $E[\varphi^{\{1-\alpha_m\}}] = 1$. The scale begins at 1 rather than 0 because log utility ($\gamma = 1$) is the boundary of the multiplicative world: admissibility is stated in terms of $E \log \varphi$, so the logarithmic agent is exactly the agent for whom the admissibility condition is the acceptance condition.

The identity that organizes the paper is immediate from (3). Because $g = \log \varphi$, the relative index is computable entirely from the additive index of the log return:

$$R_m(\varphi) = h(R(\log \varphi)), \quad h(a) = a/(1 + a), \quad (4)$$

with h strictly increasing on $(0, \infty)$. The two indices are the same theory viewed through a logarithm; every ordinal statement about additive riskinesses of log returns transfers to relative riskinesses of the corresponding multiplicative gambles, and vice versa. We use this repeatedly.

It is worth distinguishing R_m from the dimensionless quantity Aumann and Serrano themselves propose for an instrument, namely $R(\varphi - 1)$, the additive riskiness per dollar of the excess return. Both are dimensionless and both are defined on returns, but they answer different questions. The Aumann and Serrano quantity asks how risky a dollar in the instrument is and remains a dollar-normalized object built from absolute risk aversion; R_m asks how relatively risk-averse an agent must be to decline the instrument and is intrinsically a CRRA parameter. Neither is a monotone transform of the other in general. Table 1 sets the two constructions side by side.

	Aumann-Serrano (additive)	Li / Schreiber (multiplicative)
Object	g ; wealth $w \rightarrow w + g$	$\varphi > 0$; wealth $w \rightarrow w\varphi$
Admissibility	$Eg > 0, P(g < 0) > 0$	$E \log \varphi > 0, P(\varphi < 1) > 0$
Defining equation	$E[e^{\{-g/R\}}] = 1$	$E[\varphi^{\{1-\alpha_m\}}] = 1$
Calibrating class	CARA	CRRA
Critical parameter	$\alpha(g) = 1/R$, range $(0, \infty)$	$\alpha_m = 1 + \alpha(\log \varphi)$, range $(1, \infty)$
Units	Dollars	A CRRA coefficient (dimensionless)
Local risk aversion	Absolute, $\rho(w)$	Relative, $r(w)$
Monotone-RA class	DARA / IARA (Yaari 1969)	DRRA / IRRA (Li 2014a)

Table 1. Structural comparison of the additive and multiplicative indices of riskiness. The two are linked by the identity $R_m(\varphi) = h(R(\log \varphi))$, $h(a) = a/(1 + a)$.

3. Closed form, an existence characterization, and horizon invariance

This section develops the three theoretical results that make the multiplicative index a usable empirical instrument. None requires machinery beyond Aumann and Serrano and Li; the contribution is to assemble them, to phrase the existence result as a characterization with a converse, and to draw the tail asymmetry sharply. The index itself is due to Li (2014a) and Schreiber (2014), to whom any priority belongs.

3.1 The log-normal closed form

Let an asset have log return $X = \log \varphi \sim N(m, s^2)$ with $m > 0$, so φ is log-normal and admissible: $E \log \varphi = m > 0$ and $P(\varphi < 1) = P(X < 0) > 0$ for any $s > 0$.

Proposition 1 (closed form). If $X = \log \varphi \sim N(m, s^2)$ with $m > 0$ and $s > 0$, then $\alpha_m(\varphi) = 1 + 2m/s^2$ and $R_m(\varphi) = s^2/(s^2 + 2m)$. In particular $R_m(\varphi)$ exists for every such asset and depends only on (m, s^2) .

Proof. With $X \sim N(m, s^2)$, $E[e^{\beta X}] = \exp(\beta m + \beta^2 s^2/2)$. Setting this to 1 gives $\beta(\beta s^2/2 - m) = 0$, whose positive root is $\beta = 2m/s^2 = \alpha(\log \varphi)$. Substitute into (3).

The index is thus obtained from two moments with no root-finding, and each value reads as a CRRA threshold. We flag a pitfall that persists in applied work. The Aumann and Serrano “normal formula” $R = \sigma^2/2\mu$ gives the riskiness of a gamble only when *arithmetic* returns are normal. If instead *log* returns are normal, the dollar payoff is a shifted log-normal, and $\sigma^2/2\mu$ computed from log-return moments is not its riskiness; it is $R(\log \varphi) = s^2/2m$, precisely the additive object out of which α_m is built by (3). Applied work reporting Aumann and Serrano riskiness for assets specified by log-return moments should state which quantity it means. We quantify the resulting discrepancy in the panel in Section 8.

3.2 Existence: a tail characterization

The organizing claim of the paper is that the multiplicative index is the better-behaved object on financial returns. We now make that precise. The two indices read opposite tails of the return distribution, and the direction of the resulting existence asymmetry can go either way depending on which tail is heavy.

Proposition 2 (existence dichotomy). Let φ be an admissible multiplicative gamble with log return $X = \log \varphi$, and for $c > 0$ consider the additive dollar positions $g_1 = c(\varphi - 1)$ (long) and $g_2 = c(1 - \varphi)$ (short).

(i) *Left tail governs R_m .* $R_m(\varphi) \in (0, 1)$ exists if and only if $E[\varphi^{-\beta}] < \infty$ for some $\beta > 0$, equivalently, the moment generating function of X is finite on a left neighborhood of 0. This is a condition on the left tail of X (the mass of φ near 0) alone. When it fails, $\alpha(\log \varphi) = 0$ and $R_m(\varphi) = 1$ by the Schulze (2014) convention.

(ii) *Right tail governs the additive short index.* $R(g_2)$ is finite if and only if the moment generating function of φ is finite at some positive argument, a condition on the right tail of φ .

The long index $R(g_1)$ has no such obstruction ($0 < e^{\{-\beta c\varphi\}} \leq 1$ for $\beta > 0$) but has no elementary form when X is normal.

(iii) *Dichotomy.* There exist admissible φ on which R_m exists in closed form while $R(g_2) = +\infty$ (heavy right tail, light left tail, e.g. log-normal φ), and admissible φ on which $R(g_1)$ and $R(g_2)$ are finite while $R_m(\varphi) = 1$ fails to be interior (heavy left tail, light right tail).

Proof. (i) Let $f(\beta) = E[\varphi^{\{-\beta\}}] - 1$ for $\beta \geq 0$, taking the value $+\infty$ where the expectation diverges. f is convex with $f(0) = 0$ and $f'(0) = -E \log \varphi < 0$. Since $P(\varphi < 1) > 0$, on $\{\varphi < 1\}$ we have $\varphi^{\{-\beta\}} \rightarrow \infty$ as $\beta \rightarrow \infty$, so by monotone convergence $E[\varphi^{\{-\beta\}}] \rightarrow \infty$. If f is finite on some $(0, \beta_0)$ it eventually turns positive and, being convex with negative right-derivative at 0, has a unique positive root $\alpha(\log \varphi)$; then $R_m = 1/(1 + \alpha(\log \varphi)) \in (0, 1)$. The only obstruction is finiteness near 0, which involves only small φ , i.e. the left tail of X . (ii) For the short payoff the Aumann and Serrano equation is $e^{\{-\beta c\}} E[e^{\{\beta c\varphi\}}] = 1$, finite only where the mgf of φ is finite; for log-normal φ it diverges for every $\beta > 0$. For the long payoff $E[e^{\{-\beta c\varphi\}}]$ is a log-normal Laplace transform, finite but with no elementary form. (iii) Log-normal φ gives the first case by Proposition 1 and part (ii). For the second, take φ bounded above with a left tail heavy enough that $E[\varphi^{\{-\beta\}}] = \infty$ for all $\beta > 0$ (mass accumulating at 0 like a negative power); then R_m fails to be interior while both dollar positions are bounded random variables, so both additive indices are finite.

The economic reading is important. The multiplicative index depends on $E[\varphi^{\{-\beta\}}]$, the left tail of φ ; the additive short index depends on $E[e^{\{\beta\varphi\}}]$, the right tail. Gross asset returns have heavy right tails and comparatively light left tails, they are bounded below by zero under limited liability, so the empirically relevant case is exactly the one in which the multiplicative index is finite and closed-form while the additive short index is infinite. This is the case Li (2014a) flags in a footnote, and that Schulze (2014) extends the additive index to accommodate; Proposition 2 shows it is not a special embarrassment of one example but the generic situation for returns, and it identifies the converse regime in which the additive index is the survivor. The inversion of the usual presumption, additive index primitive, multiplicative index derivative, is thereby a statement about tails, provable in closed form.

3.3 Products, dependence, and horizon invariance

Multi-period and portfolio questions concern products $\prod_i \varphi_i$, and $\log \prod_i \varphi_i = \sum_i \log \varphi_i$ turns them into the sums studied by Aumann and Serrano (2008) and Li (2014b). Because $R_m(\varphi) = h(R(\log \varphi))$ with h strictly increasing, order relations among additive riskinesses of log returns carry over to relative riskinesses of the corresponding products.

Proposition 3 (products and horizon). Let $\varphi_1, \dots, \varphi_n$ be multiplicative gambles. (i) If they are i.i.d., then $R_m(\prod_i \varphi_i) = R_m(\varphi_1)$ for every n . (ii) If $\log \varphi_1$ and $\log \varphi_2$ are positively quadrant dependent (Lehmann, 1966), then $R_m(\varphi_1 \varphi_2) \geq \min\{R_m(\varphi_1), R_m(\varphi_2)\}$; if negatively quadrant dependent, then $R_m(\varphi_1 \varphi_2) \leq \max\{R_m(\varphi_1), R_m(\varphi_2)\}$.

Proof. (i) Aumann and Serrano show the additive riskiness of a sum of i.i.d. gambles equals that of a single summand; apply h to $R(\sum_i \log \varphi_i) = R(\log \varphi_1)$. (ii) Li (2014b) shows $R(g_1 + g_2) \geq \min\{R(g_1), R(g_2)\}$ under positive quadrant dependence and the reverse under negative dependence; apply h with $g_i = \log \varphi_i$. Quadrant dependence is preserved by the increasing map $\log$, so the hypotheses may be stated on (φ_1, φ_2) directly.

Corollary (horizon invariance). For an i.i.d.-return asset the relative index of the buy-and-hold position is the same at every horizon. Under log-normal returns this is immediate from Proposition 1: the n -period log return is $N(nm, ns^2)$, so $R_m = ns^2/(ns^2 + 2nm) = s^2/(s^2 + 2m)$.

Economically, positive dependence floors the compounded riskiness by the less risky factor and negative dependence caps it by the riskier, the diversification reading, now for products rather than sums. When returns are *not* i.i.d., horizon invariance breaks in a way we can measure. Writing the T -horizon critical coefficient in terms of the Lo-MacKinlay (1988) variance ratio $VR(T)$,

$$\alpha_m(T) - 1 = (\alpha_m(1) - 1) / VR(T), \quad (5)$$

so that mean reversion ($VR < 1$) raises the critical coefficient with the horizon and momentum or persistence ($VR > 1$) lowers it. We use (5) in Section 8 to show that a rolling Treasury-bill position and the market portfolio, far apart at a one-year horizon, converge at a ten-year horizon because bill returns are strongly persistent.

4. Related literature

The literature on objective riskiness has developed along the additive index; we place our contribution against it briefly. The bibliographic record should be stated once: the theory used here is Li (2014a), *Economic Theory* 55(2), with the dependence results split into Li (2014b); the relative index was obtained independently and nearly simultaneously by Schreiber (2014), *Economic Theory* 56(2), who frames the paper as being about the relative index and has become the more-cited reference. Any work in this area should cite both, and we report Li's α_m while adopting Schreiber's motivating observation that an asset can be riskier in relative terms while less risky in absolute terms.

Existence and computation. Schulze (2014) characterizes which gambles admit the additive index, extends it to the rest via a duality-consistent construction, and methodologically closest to our Section 3, connects the problem to the bilateral Laplace transform, obtaining closed forms for the exponential, Gamma and Variance-Gamma cases and showing the riskiness of a Lévy process is time-invariant. Homm and Pigorsch (2012b) treat existence and the normal-inverse-Gaussian case; Cui, Wu and Zhu (2023) give an explicit solution. Our Proposition 2 is the multiplicative analogue: it locates existence in the left tail of the log return and pairs it with a converse.

Axiomatics, nesting and limits. Foster and Hart (2009) give the other operational measure of riskiness; Bali, Cakici and Chabi-Yo (2011) nest both and add an options-implied version; Hart (2011) supplies the wealth- and utility-uniform dominance orders. Heller and Schreiber (2020) and Schreiber (2016) show that for continuous processes at short horizons the Aumann and Serrano, Foster-Hart and Schreiber measures collapse to a common variance-to-mean index, a short-horizon limit complementary to our i.i.d. horizon invariance, which is exact rather than asymptotic.

Applications, and the empirical gap. The applied literature is real but additive. Homm and Pigorsch (2012a) and Kadan and Liu (2014) use the Aumann and Serrano index as a Sharpe ratio robust to non-normality; further applications span hedging (Chen, Ho and Tzeng, 2014; Ehsani and Lien, 2015), the equity premium (Bali, Cakici and Chabi-Yo, 2015), executive compensation

(Lu, Shen, Shih and Tsai, 2023), cryptocurrencies (Yamawake, Sheely, Serrano and Hodoshima, 2024) and temporal aggregation (Hodoshima and Yamawake, 2020, 2022). Against this, the multiplicative index has one empirical application: Lu, Yang and Wong (2018), in the *Annals of Financial Economics*, apply Schreiber's relative index to test time diversification on the S&P 500, S&P 400 and NASDAQ. That is a time-series study of three indices. No cross-sectional application exists, and none has used the index to characterize the size and value dimensions along which the empirical asset-pricing literature is organized. That is the gap this paper fills.

5. Data and research design

We assemble annual log returns from 1928 to 2025 for 27 assets and partition the 98-year sample into five non-overlapping periods of roughly two decades, 1928-1945, 1946-1965, 1966-1985, 1986-2005 and 2006-2025, so that each period carries enough observations to estimate a mean and variance while preserving genuine regime variation across the century. The data are from publicly available websites (French, 2026 and Damodaran, 2026). The cross-section is built to isolate the two canonical sorts of the empirical asset-pricing literature. Ten *size deciles* (Size_Dec_1 smallest to Size_Dec_10 largest) span the market-capitalization sort; ten *book-to-market deciles* (Value_Dec_1 growth to Value_Dec_10 deep value) span the value sort; and seven aggregate classes, the aggregate market, the S&P 500, three-month Treasury bills, ten-year Treasury bonds, Baa corporate bonds, real estate and gold, anchor the equity deciles against fixed income and real assets. With five periods this yields a panel of $27 \times 5 = 135$ asset-period cells and, for pairwise rank analysis, 1,755 within-period asset pairs.

For each cell we take the working assumption of the theory that annual log returns are normal and price each risky asset at \$100. The multiplicative index is then computed in closed form from Proposition 1, $\alpha_m = 1 + 2m/s^2$, and the additive index of the corresponding \$100 long position is computed numerically: the Aumann and Serrano equation $E[\exp\{-\beta \cdot 100(e^{\{X\}} - 1)\}] = 1$ is solved for β by 100-node Gauss-Hermite quadrature of the log-normal expectation combined with Brent root-finding, cross-checked against Monte Carlo. The additive index of the short position is recorded as non-existent, per Proposition 2(ii). Section 8 reports a non-parametric check that does not impose normality, replacing the closed form with the empirical moment generating function of the realized log returns.

The pricing test of Section 9 requires a higher-frequency panel and a matched set of test assets. For it we use monthly value-weighted returns from the Kenneth French data library over 1926:07-2026:05: the ten size deciles and the ten book-to-market deciles that correspond to our annual sorts, together with the market factor and the one-month Treasury bill. These returns are reported as simple percentages; to keep the riskiness predictor on the same footing as the annual panel we convert them to log returns, $\log(1 + r/100)$, and we verify that the annual log-return panel is exactly the calendar-year sum of these monthly log returns (to five decimal places). The choice of return space is not incidental. The multiplicative index is a functional of log returns, so its predictor must be built from them; a tradable portfolio and a linear pricing slope, by contrast, aggregate across assets, where simple returns add and log returns do not. We therefore build the predictor from log returns and take the priced quantity to be the simple excess return, and we confirm in Section 9 that using a log-excess left-hand side instead leaves every conclusion intact.

6. Existence and ordinal agreement

Two facts frame everything that follows. First, the existence pattern of Proposition 2 holds in every one of the 135 cells: the multiplicative index is available in closed form from two moments, the additive index of the long position requires numerical solution of an equation with no elementary form, and the additive index of the short position does not exist. The theory's central asymmetry is not a knife-edge; it is the universal state of the panel.

Second, at the level of a full-panel ranking the two indices agree almost perfectly. Within each period the Spearman rank correlation between α_m and the additive index of the long position lies between 0.9951 and 0.9994, and across all 1,755 within-period pairs only 12, 0.68%, are discordant. An investor who ranked the cross-section by the additive dollar index and an investor who ranked it by the critical CRRA coefficient would order all but one pair in seventy the same way. This is not a surprise given (4): R_m is a monotone transform of $R(\log \varphi)$, and $R(\log \varphi)$ and $R(\varphi - 1)$ are close when the variance drag separating log and arithmetic moments is small.

But the two indices are not identical, and the disagreements are informative. The 12 discordant pairs are not scattered across the panel; they concentrate among the value deciles, consisting of value deciles: an asset that the additive dollar index rates as the riskier of a pair is, on the CRRA scale, the safer, because its higher mean log return buys down the relative index by more than its higher variance raises it. Schreiber's motivating claim that an asset may be riskier in relative terms while less risky in absolute terms is therefore not merely a theoretical possibility; it occurs in long-run U.S. data, and it occurs specifically along the value dimension. That localization is the first sign of the structural result in Section 7.

7. The non-monotonicity of the critical CRRA coefficient across value

The paper's central finding concerns the shape of α_m across the book-to-market sort. A dollar-denominated risk measure records the value book as a monotone gradient in mean and volatility; the CRRA-calibrated index does not. Averaged over the five periods with period fixed effects, α_m rises from growth to the middle of the book and falls again into deep value, tracing an interior hump. The two extremes of the value book, Value_Dec_1 and Value_Dec_10, carry a critical CRRA coefficient close to 5.6, while the interior peaks near 9.8; the extremes sit 32.6% below the interior, and the contrast is statistically strong ($t = -3.33$, $p = 0.006$).

A quadratic fit $\alpha_m = a + b \cdot d + c \cdot d^2$ in the decile index d , with period fixed effects, makes the shape precise and separates the two books cleanly. On the value book the fitted vertex falls at decile 4.88, an *interior* maximum, the hump. On the size book the same fit places the vertex at 10.01, on the boundary: α_m is (essentially) monotone across size, exactly as intuition expects, and the curvature that produces an interior optimum is specific to the value sort. Table 2 summarizes.

Sort	Vertex decile	Extreme vs. interior	t	p	Shape
Book-to-market (value)	4.88 (interior)	-32.6%	-3.33	0.006	Interior hump
Market cap (size)	10.01 (boundary)	—	—	—	Monotone

Table 2. Shape of the critical CRRA coefficient across the two equity sorts. Quadratic fit with period fixed effects; the value book has an interior maximum (a hump), the size book a boundary vertex (monotone). The value extremes sit 32.6% below the value interior.

The economics is the point, and it is legible only because α_m is calibrated in relative risk aversion. Recall $\alpha_m - 1 = 2m/s^2$: the critical coefficient is high when the mean log return is high relative to variance, and low otherwise. The two ends of the value book are unattractive to a moderately risk-averse CRRA agent for *opposite* reasons. Growth stocks (Value_Dec_1) have a low mean log return; they are unattractive for the ordinary reason, a small numerator. Deep-value stocks (Value_Dec_10) have the *highest* mean log return in the book, the raw value premium, but their variance is large enough that the variance drag $s^2/2$ pulls $2m/s^2$ back down to the level of growth. On the arithmetic scale deep value looks like a high-return asset; on the log scale that decides multiplicative acceptance, its risk-adjusted mean is no better than growth's. The hump is the geometric signature of the value premium being eaten by volatility at the extreme of the book.

This also disposes of a claim that does not generalize. On the size sort the multiplicative index *compresses* the cross-section relative to the additive one, because in a size sort the mean and variance covary in a way that partially cancels in the ratio $2m/s^2$. That compression is a feature of size sorts specifically, not of equity in general: on the value book the cross-section of α_m spreads 1.31 to 1.81 times *wider* than the corresponding additive spread in every period. The instrument that looks like it merely shrinks distances on one sort magnifies them on another, and the difference is exactly the covariation of mean and variance that the CRRA calibration is designed to weigh. Because the additive dollar index is nearly ordinally equivalent to the multiplicative one (Section 6), it traces the same hump; what a dollar-denominated measure cannot supply is the reading of that hump in units of risk aversion, and it resolves the value book with visibly less dynamic range.

The hump is identified from the shape across deciles pooled over periods, not from any single cell, and it is robust to sampling error in the estimated moments. A bootstrap that resamples the annual observations within each period two thousand times, jointly across assets, so the cross-sectional dependence is preserved, and recomputes the entire panel each time places the fitted value vertex at an interior decile in 99.7% of replications, with a 95% interval of [3.2, 6.3]; the extremes-versus-interior gap is negative in every replication (95% interval [-46%, -19%]). The size book, by contrast, yields an unstable vertex (95% interval [6.7, 37]) and a gap whose interval covers zero. Individual per-cell α_m estimates are imprecise, each rests on roughly two decades of annual data, but the *shape* is sharply identified, and the value-versus-size contrast survives the propagation of that estimation error.

8. Calibration, dependence, horizon, and robustness

8.1 The CRRA band

Because every α_m is a risk aversion threshold, the panel can be read against the range of relative risk aversion the calibration literature regards as plausible, roughly one to ten, with most estimates in the low single digits. An asset with α_m below that band is one that essentially no CRRA agent would hold; an asset with α_m above it is one essentially no CRRA agent would decline. The panel minimum critical coefficient is 1.25, and the lowest values cluster in the most volatile equity cells; in the Great Depression sub-period (1928-1945) the size book bottoms at

1.65. Values this close to 1 say that in the most turbulent regimes the riskiest equity deciles were acceptable only to agents barely more risk-averse than logarithmic, a sharp, interpretable statement that the additive dollar index, denominated in dollars rather than in risk aversion units, cannot make directly.

The conflation hazard flagged in Section 3.1 is quantified here. Reporting $\sigma^2/2m$ from log-return moments as though it were the additive dollar riskiness overstates the true per-dollar quantity by between 1.3% and a factor of 4.88 across the equity cells, with a mean overstatement near 40%. The error is largest exactly where variance drag is largest, the high-volatility deciles that drive the value hump, so the conflation would distort precisely the cells that matter most for the paper's finding.

8.2 Dependence corollaries

Proposition 3 and its relatives yield three testable bounds, and all three hold without a single violation across the panel. First, a harmonic-mean floor for the critical coefficient of an equally weighted geometric (buy-and-hold-rebalanced) portfolio, which we believe is new: the portfolio's α_m is bounded below by the harmonic mean of the constituents' coefficients, so diversification cannot make a geometric portfolio more relatively risky than its least risky component allows. Second, the sign-of-dependence bounds of Proposition 3(ii), inherited from Li (2014b): positive quadrant dependence of log returns floors the product's coefficient by the smaller and negative dependence caps it by the larger. Third, the horizon law (5). Across all within-period pairs, every pairwise product bound holds; the geometric six-asset diversification ratio reaches 17.09, and the numerical convergence of the closed form against the quadrature solution is at the level of 10^{-14} , confirming the computation.

8.3 Horizon

Equation (5) turns the panel's autocorrelation structure into horizon-dependent critical coefficients. The Lo-MacKinlay variance ratios reveal strong persistence in rolling Treasury-bill returns, $VR(10) \approx 9.17$, a z -statistic against the i.i.d. null above 15, so a rolling-bill position that is far safer than the market portfolio at a one-year horizon converges toward it at a ten-year horizon, as the persistence inflates the effective variance of the compounded bill position. The market portfolio's own variance ratio is near one, so its critical coefficient is essentially horizon-free, consistent with the i.i.d. corollary. Horizon invariance is thus not a universal fact but a diagnostic: departures from it locate exactly the assets whose serial dependence matters for long-horizon risk.

8.4 Robustness to the normality assumption

The closed form of Proposition 1 assumes normal log returns, and that assumption is not innocuous here: a Shapiro-Wilk test rejects normality at the 5% level in 21 of the 135 cells (16%), including 8 of the 50 value-decile cells. Because the paper's finding is a shape result in the critical coefficient, it must not be an artifact of the log-normal model. The natural check replaces the closed form with a non-parametric estimate. For each cell, form the empirical moment generating function of the realized log returns, $\hat{w}(\beta) = n^{-1} \sum_t e^{\{-\beta X_{it}\}}$, and solve $\hat{w}(\beta) = 1$ for the positive root $\beta = \alpha(\log \varphi)$, giving $\hat{\alpha}_m = 1 + \beta$ directly from the data with no distributional assumption. This is the multiplicative analogue of the non-parametric Aumann and Serrano estimator of Homm and Pigorsch (2012a), and it makes Proposition 2 concrete: a positive root

exists if and only if the realized sample contains at least one down year, and in exactly the seven near-riskless cells with no in-sample loss (Treasury bills and real estate in particular periods) the estimator is $+\infty$, the left-tail existence condition of Proposition 2(i) in observable form.

Across the 128 cells with a finite root the two estimators agree where it matters. The within-period rank correlation between $\hat{\alpha}_m$ and the closed-form α_m ranges from 0.923 to 0.997; the median absolute deviation in the critical coefficient is 0.70 and the median relative deviation is +0.3%. The few large discrepancies are confined to the low-variance, near-riskless cells where α_m is very large and its exact level is ordinally irrelevant; on the equity deciles that generate the value hump the two estimators are close.

Most important, the shape result is not an artifact of normality. Re-running the value/size shape test on the non-parametric $\hat{\alpha}_m$ (Table 3) preserves every qualitative feature and sharpens the value result. The value book's fitted vertex moves from 4.88 to 5.01, still an interior maximum; the size book's vertex stays on the boundary (10.01 $\rightarrow$ 10.10, monotone); and the extremes-versus-interior gap on the value book widens from -32.6% to -37.8% , with the contrast becoming *more* significant ($p = 0.006 \rightarrow 0.001$), while the size book stays flat and insignificant. The hump is a fact about U.S. equity returns, not about the log-normal model used to summarize them.

	Value book		Size book	
Estimator $\rightarrow$	Closed form	Non-param.	Closed form	Non-param.
Vertex decile	4.88	5.01	10.01	10.10
Shape	interior hump	interior hump	boundary	boundary
Extremes vs. interior	-32.6%	-37.8%	-10.7%	-8.4%
t (extremes)	-3.33	-3.87	-0.97	-0.85
p (extremes)	0.006	0.001	0.357	0.416
t (quadratic k^2)	-2.94	-3.19	-2.46	-1.78

Table 3. Robustness of the shape result to the normality assumption. The value/size quadratic fit (period fixed effects) is re-estimated on the non-parametric critical coefficient $\hat{\alpha}_m$ obtained from the empirical moment generating function of realized log returns, alongside the closed-form estimate. The interior hump on the value book and the boundary (monotone) vertex on the size book are preserved; the value contrast sharpens. Estimated on the 128 of 135 cells with an in-sample down year.

8.5 Relation to the Sharpe ratio

A natural objection is that the multiplicative index is merely a relabelling of the Sharpe ratio, and the objection deserves a direct answer. It is partly right and instructively wrong. Writing the log Sharpe ratio as $SR = m/s$, the identity $\alpha_m - 1 = 2m/s^2 = 2 \cdot SR/s$ shows that α_m is a *volatility-tilted* Sharpe ratio: two assets with the same Sharpe ratio but different volatilities carry different critical coefficients, the lower-volatility asset the higher one. The two are therefore not ordinally equivalent. Across the panel the Spearman correlation between α_m and the log Sharpe ratio ranges from 0.71 to 0.97 by period, high, but bounded well away from unity, and within the value book the two rankings disagree on four to eight of the forty-five pairs in every period.

What the Sharpe ratio does *not* do is dispose of the hump. Re-running the value-book shape test on the log Sharpe ratio itself returns an interior vertex (decile 5.2) and a negative extremes-versus-interior gap (-29% , $p = 0.023$), close to but distinct from the α_m result (decile 4.9, -33% , $p = 0.006$). The hump is thus not an artifact of the CRRA calibration specifically; it is a robust feature of the risk-adjusted returns of the value book, visible to any mean-over-volatility measure. The contribution of α_m is not exclusive detection but interpretation: because its units are a relative risk aversion coefficient, the hump reads directly as a statement about *which agents hold which deciles*, the interior of the book is acceptable to CRRA agents up to $\gamma \approx 9$, both extremes only up to $\gamma \approx 5.6$, and because it is a genuinely different ranking from the Sharpe ratio, the two can and do order the value book differently. The index earns its place as an interpretable, volatility-sensitive complement to the Sharpe ratio, not as a substitute that sees something no other measure can. Whether either the hump or the ranking it refines is priced in the cross-section is the question we turn to next.

9. Does relative riskiness price the cross-section?

The preceding sections establish the critical CRRA coefficient as a coherent descriptive instrument and locate a robust non-monotonicity in it across the value book. Whether that structure is *priced*, whether assets the index rates as riskier earn higher average returns, and whether the value hump leaves a trace in expected returns, is a separate question, and the annual panel is poorly suited to answer it. Five non-overlapping periods furnish only a handful of cross-sections; a slope estimated on them carries almost no power, and an earlier annual attempt returned a point estimate near zero that could not distinguish “no relation” from “no power.” We therefore take the question to monthly data, where roughly a thousand cross-sections make a properly powered test possible.

The design is a standard predictive cross-sectional test with the multiplicative index as the sorting characteristic. At each month-end, for each of the twenty size and value deciles, we estimate the relative index R_m over a trailing window of past monthly log returns $[t-w, t-1]$, admissible ($m > 0$, $s > 0$) in 89% of asset-months, and relate the lagged R_m to the decile’s return in month t . Because the index is horizon-invariant for i.i.d. returns (Proposition 3), the trailing R_m sits on the same scale as the annual estimates used above, and it recovers the same structure: averaged over the sample the small-cap deciles carry higher R_m than large, and the value book traces the familiar interior trough in R_m , the α_m hump, with the riskiest readings at the growth and deep-value extremes. The predictor is built from log returns; the priced quantity is the simple excess return $r_i - r_{\text{f}}$; and we report a 60-month window as the primary specification, with a 120-month window and a log-excess left-hand side as robustness. Four tests follow: the first three ask whether the *level* of relative riskiness predicts returns, and the fourth whether *exposure* to a riskiness factor is priced.

Fama-MacBeth. Regressing each month’s cross-section of excess returns on the lagged, cross-sectionally standardized R_m and averaging the slopes (Newey-West standard errors, six lags) yields a slope indistinguishable from zero: across all twenty deciles the average slope is 0.001% per month ($t = 0.03$) at the 60-month window and 0.035% ($t = 1.03$) at 120 months. The null holds in every cut, size deciles alone, value deciles alone, the matched 1928-2025 window, and a log-excess left-hand side. The single positive flicker, the size deciles at the 120-month window ($t = 2.03$ under a naïve standard error), is the size premium re-expressed through the predictor rather than riskiness pricing: it vanishes under the log-excess left-hand side ($t = 0.09$) and at the

60-month window ($t = 0.81$). The near-zero annual slope was not an artifact of low power; with an order of magnitude more cross-sections the slope stays flat.

A riskiness-sorted portfolio. Sorting the twenty deciles each month into quintiles by trailing R_m and holding the riskiest quintile against the safest, an equal-weighted, monthly-rebalanced long-short, produces a spread of -0.02% per month ($t = -0.15$), with mean excess returns essentially flat across the five quintiles (0.82, 0.79, 0.83, 0.83, 0.81% per month from safe to risky). The spread's CAPM alpha is small and, if anything, negative (-0.19% per month, $t \approx -1.9$), a faint low-volatility tilt rather than the positive premium a priced riskiness factor would produce. The size-only spread is monotone in R_m (rank correlation $+0.90$) but insignificant ($t = 0.50$) and carries no CAPM alpha, the size premium, absorbed by the market factor.

The value hump, forward. An earlier annual exercise found a suggestive positive rank correlation ($+0.63$, $n = 40$) between a decile's α_m and its subsequent value-book return which, taken at face value, would have hinted that the hump lines up with expected returns. It does not survive. The monthly analogue, the average cross-sectional rank correlation between trailing α_m and the next-month value-decile return, is $+0.02$ ($t = 1.5$), and the forward returns of the value book are not humped but monotone increasing in book-to-market (0.64% per month at growth rising to 0.90% at deep value), which is simply the value premium. The α_m hump is a statement about risk-adjusted mean log returns, the value premium net of variance drag, not about which deciles subsequently earn more. The annual association was a small-sample artifact of four period transitions.

Test	Cut / window	Estimate	t (NW)	Reading
Fama-MacBeth slope	all 20, 60m	0.001%/mo	0.03	no pricing
Fama-MacBeth slope	all 20, 120m	0.035%/mo	1.03	no pricing
Fama-MacBeth slope	size 10, 120m	0.079%/mo	1.81	size premium; dies under log-LHS (0.09) and 60m (0.81)
Fama-MacBeth slope	value 10, 60m	-0.028% /mo	-0.77	no pricing
Riskiest-safest spread	all 20, 60m	-0.02% /mo	-0.15	no premium
Spread CAPM alpha	all 20, 60m	-0.19% /mo	-1.9	faint low-vol tilt
Rank corr(α_m , fwd)	value book	$+0.02$	1.5	annual $+0.63$ does not survive

Table 4. Monthly asset-pricing tests of the multiplicative index. Twenty size and value deciles, monthly value-weighted returns 1926:07-2026:05; predictor is the trailing-window relative index R_m built from log returns (60-month primary, 120-month robustness); Newey-West standard errors with six lags. Slopes are per-month excess return per cross-sectional standard deviation of R_m . No specification delivers a positive, significant riskiness premium.

A stochastic-discount-factor test. The three tests above ask whether the *level* of relative riskiness predicts returns; the covariance-side question is whether *exposure* to a riskiness factor is priced, and it is answered by a stochastic-discount-factor test rather than a characteristic regression. We take the monthly riskiest-minus-safest spread as a candidate factor f and estimate the linear SDF restriction $E[m \cdot R^e] = 0$, with $m = 1 - b(f - E[f])$, across the twenty deciles by two-step GMM with a Newey-West weighting matrix, reporting the loading b , the implied price of risk $\lambda = b \cdot \text{Var}(f)$, Hansen's J -test of the nineteen over-identifying restrictions, and the Hansen-Jagannathan distance (Table 5). The riskiness factor is not priced: its loading ($b = 1.9$, $t = 1.1$) and its price of risk ($\lambda = 0.21\%$ per month, $t = 1.1$) are both statistically zero, and the point estimate of b is unstable across weighting matrices (a first-stage +8.4 collapsing to +1.9 under the efficient weight), the signature of an irrelevant factor. The J -statistic rejects (57.1, $p < 0.001$), but this must be read in context. Run through the identical machinery, the market excess return is *also* rejected (34.1, $p = 0.02$): no single factor spans the size and value cross-section, so a J -rejection is the norm on these test assets rather than a verdict specific to riskiness. The market benchmark instead validates the procedure and sets the scale. Its loading is strongly significant ($t = 4.5$), and its implied price of risk, 0.72% per month, essentially reproduces its own mean of 0.69%, as a traded factor that prices itself must, so the machinery detects a premium when one is present. Against that scale the riskiness factor leaves a *larger* Hansen-Jagannathan distance (0.22 versus 0.18) than the market: it is a worse pricing kernel than a single-factor CAPM, not a better one. The discriminating quantity is not the J -test, which rejects broadly, but the price of risk, which is zero for riskiness and large for the market.

Factor	b (t)	λ %/mo (t)	$E[f]$ %/mo	J (df 19)	p	HJ dist.
Riskiness spread (R_m Q5–Q1)	1.9 (1.1)	0.21 (1.1)	−0.02	57.1	<0.001	0.22
Market (Mkt–RF, benchmark)	3.6 (4.5)	0.72 (4.5)	0.69	34.1	0.018	0.18

Table 5. GMM stochastic-discount-factor tests. The SDF is $m = 1 - b(f - E[f])$; test assets are the twenty size and value deciles (monthly simple excess returns, 1931:07-2026:05); two-step GMM with a Newey-West optimal weighting matrix (six lags); J tests the nineteen over-identifying restrictions. λ is the implied price of factor risk; for a traded factor a correctly specified model implies $\lambda = E[f]$. The riskiness factor carries a price of risk indistinguishable from zero and prices the cross-section worse than the market; the market benchmark, run through the identical machinery, is strongly priced and reproduces its own mean, confirming that the procedure detects a premium when one is present.

Read together, the four tests deliver a clean and well-powered null, and they do so from both directions: the characteristic tests find no predictive slope on the level of riskiness, and the stochastic-discount-factor test finds no price on exposure to it. Relative riskiness orders the cross-section sensibly and exhibits a genuine value-dimension hump, but that ordering is not a separately priced risk factor: the average returns of these deciles are governed by the size and value premia, which the multiplicative index does not encode, and the hump does not translate into a return pattern. This is not a failure of the instrument but a delineation of what it is for. The

objective riskiness indices, additive or multiplicative, measure the left-tail structure of the return distribution, a quantity of direct use for portfolio screening, performance evaluation, and the acceptance-threshold reading developed in Section 8; they are not, on this evidence, cross-sectional pricing factors. Reporting the null also forecloses an over-reading the annual panel alone would have invited: that the value hump is a priced anomaly. It is not.

10. Conclusion

Finance deals in multiplicative gambles, and the objective index built for them, calibrated through relative rather than absolute risk aversion, is both better behaved on financial returns and, until now, essentially unused in the cross-section. This paper makes it usable and uses it. On the theory side, the index has a two-moment closed form under lognormality, and its existence is governed by the left tail of the log return, opposite the right tail that governs the additive short index; the resulting dichotomy makes precise, and gives a converse to, the intuition that the multiplicative index is the natural object for returns. Through the logarithm the index is horizon-invariant for i.i.d. returns and inherits dependence bounds as bounds on products.

On the evidence side, the two indices agree almost perfectly as panel rankings, but where they disagree they disagree along the value dimension, and the reason is structural: the critical CRRA coefficient is non-monotone across the book-to-market sort, humped in the interior, while it is monotone across size. Deep value carries the highest mean log return in the book yet the same low critical coefficient as growth, because its variance drag consumes the premium, a statement about the shape of the value premium's riskiness that is available only on the scale that decides multiplicative acceptance. The compression the index produces on a size sort does not generalize; on a value sort it resolves the cross-section more finely, not less, and the hump survives non-parametric estimation from the empirical moment generating function, so it is a fact about the returns rather than about the log-normal model. The instrument the empirical literature left on the shelf agrees with the additive index on the ranking, yet reads it in units, relative risk aversion, that make the shape of the value premium's riskiness interpretable, and it rests on theory the additive index lacks.

The evidence side carries a second, negative result that bounds the first. With monthly data and roughly a thousand cross-sections we test whether the riskiness ordering, and the value hump within it, is priced, and it is not: the Fama-MacBeth slope of return on trailing relative riskiness is flat, a riskiest-minus-safest decile portfolio earns nothing and carries no price of risk in a stochastic-discount-factor test, and the annual hint that the hump aligned with subsequent value returns dissolves under power. The hump is real and robust, but it is a property of the return distribution rather than a priced anomaly; average returns across these deciles answer to the size and value premia, which the index does not encode. Far from weakening the paper, the null sharpens its claim: the multiplicative index is an interpretable instrument for measuring riskiness and screening portfolios, not a factor that prices the cross-section, and saying so plainly keeps the descriptive contribution from being mistaken for a pricing one.

The relative index itself is not ours; it is due to Li (2014a) and Schreiber (2014), and the contribution here is the existence characterization, the products/horizon results, and the cross-sectional evidence. And individual ranking reversals in century-long return series are delicate, small differences in mean and variance can flip a single pair, so the findings that carry the paper are the panel-wide shapes, which are robust and survive non-parametric estimation. Natural

extensions include the double sort on size and value and an options-implied version of the multiplicative index in the spirit of Bali, Cakici and Chabi-Yo (2011). The cross-sectional pricing question is settled here in the negative; the natural sequel is whether the index earns its keep where it is built to work, as an out-of-sample screening and performance instrument, rather than as a pricing factor.

References

- Arrow, K. J., 1965. Aspects of the Theory of Risk-Bearing. Yrjö Jahnssonin Säätiö, Helsinki.
- Artzner, P., Delbaen, F., Eber, J.-M., Heath, D., 1999. Coherent measures of risk. *Mathematical Finance* 9 (3), 203-228.
- Aumann, R. J., Serrano, R., 2008. An economic index of riskiness. *Journal of Political Economy* 116 (5), 810-836.
- Bali, T. G., Cakici, N., Chabi-Yo, F., 2011. A generalized measure of riskiness. *Management Science* 57 (8), 1406-1423.
- Bali, T. G., Cakici, N., Chabi-Yo, F., 2015. A new approach to measuring riskiness in the equity market: Implications for the risk premium. *Journal of Banking & Finance* 57, 101-117.
- Chen, Y.-T., Ho, K.-Y., Tzeng, L. Y., 2014. Riskiness-minimizing spot-futures hedge ratio. *Journal of Banking & Finance* 40, 154-164.
- Cui, Z., Wu, C., Zhu, L., 2023. Explicit solution to the economic index of riskiness. *Economics Letters* 232, 111328.
- Damodaran, A., 2026. Annual returns on investments in the United States, 1928-2024. NYU Stern School of Business. Retrieved from <https://pages.stern.nyu.edu/~adamodar/>
- Dybvig, P. H., Lippman, S. A., 1983. An alternative characterization of decreasing absolute risk aversion. *Econometrica* 51 (1), 223-224.
- Ehsani, S., Lien, D., 2015. A note on minimum riskiness hedge ratio. *Finance Research Letters* 15, 11-17.
- Foster, D. P., Hart, S., 2009. An operational measure of riskiness. *Journal of Political Economy* 117 (5), 785-814.
- French, K. R., 2026. Kenneth R. French Data Library. Dartmouth College Tuck School of Business. Retrieved from https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html
- Hart, S., 2011. Comparing risks by acceptance and rejection. *Journal of Political Economy* 119 (4), 617-638.
- Heller, Y., Schreiber, A., 2020. Short-term investments and indices of risk. *Theoretical Economics* 15 (3), 1247-1278.
- Hodoshima, J., Yamawake, T., 2020. The Aumann-Serrano performance index for multi-period gambles in stock data. *Journal of Risk and Financial Management* 13 (11), 288.

Hodoshima, J., Yamawake, T., 2022. Temporal aggregation of the Aumann-Serrano and Foster-Hart performance indexes. *International Review of Financial Analysis* 83, 102282.

Homm, U., Pigorsch, C., 2012a. Beyond the Sharpe ratio: An application of the Aumann-Serrano index to performance measurement. *Journal of Banking & Finance* 36 (8), 2274-2284.

Homm, U., Pigorsch, C., 2012b. An operational interpretation and existence of the Aumann-Serrano index of riskiness. *Economics Letters* 114 (3), 265-267.

Kadan, O., Liu, F., 2014. Performance evaluation with high moments and disaster risk. *Journal of Financial Economics* 113 (1), 131-155.

Lehmann, E. L., 1966. Some concepts of dependence. *Annals of Mathematical Statistics* 37 (5), 1137-1153.

Li, M., 2014a. On Aumann and Serrano's economic index of risk. *Economic Theory* 55 (2), 415-437.

Li, M., 2014b. Aumann and Serrano's economic index of risk for sums of gambles. *Cogent Economics & Finance* 2 (1), 921574.

Lo, A. W., MacKinlay, A. C., 1988. Stock market prices do not follow random walks: Evidence from a simple specification test. *Review of Financial Studies* 1 (1), 41-66.

Lu, C.-C., Shen, C. H.-h., Shih, P.-T., Tsai, W.-C., 2023. Option implied riskiness and risk-taking incentives of executive compensation. *Review of Quantitative Finance and Accounting* 60 (3), 1143-1160.

Lu, R., Yang, C., Wong, W.-K., 2018. Time diversification: Perspectives from the economic index of riskiness. *Annals of Financial Economics* 13 (3), 1850012.

Palacios-Huerta, I., Serrano, R., 2006. Rejecting small gambles under expected utility. *Economics Letters* 91 (2), 250-259.

Pratt, J. W., 1964. Risk aversion in the small and in the large. *Econometrica* 32 (1-2), 122-136.

Rabin, M., 2000. Risk aversion and expected utility: A calibration theorem. *Econometrica* 68 (5), 1281-1292.

Schreiber, A., 2014. Economic indices of absolute and relative riskiness. *Economic Theory* 56 (2), 309-331.

Schreiber, A., 2016. Comparing local risks by acceptance and rejection. *Mathematical Finance* 26 (2), 412-430.

Schulze, K., 2014. Existence and computation of the Aumann-Serrano index of riskiness and its extension. *Journal of Mathematical Economics* 50, 219-224.

Yaari, M. E., 1969. Some remarks on measures of risk aversion and on their uses. *Journal of Economic Theory* 1 (3), 315-329.

Yamawake, T., Sheely, J., Serrano, R., Hodoshima, J., 2024. Comparative performance of cryptocurrencies through the Aumann-Serrano economic index of riskiness. *Annals of Operations Research*, forthcoming.

Appendix A. Illustrative closed-form values

Table A1 applies the closed form (1) to five stylized asset classes to illustrate how each α_m reads as a CRRA threshold and how variance drag separates the log mean from the arithmetic mean. The moments are illustrative, not panel estimates. A CRRA agent holds the asset if and only if $\gamma < \alpha_m$.

Asset class	m (%)	s (%)	Arith. mean (%)	R_m	α_m (crit. γ)
U.S. equity market	6.5	18.0	8.46	0.200	5.01
Small-cap equity	8.5	30.0	13.88	0.346	2.89
Corporate (Baa) bonds	3.2	8.5	3.63	0.101	9.86
Long Treasuries	2.2	9.0	2.64	0.155	6.43
Gold	2.8	16.5	4.25	0.327	3.06

Table A1. Illustrative annual log-return moments (illustrative only). R_m and α_m from equation (1); arithmetic mean $\exp(m + s^2/2) - 1$. Small-cap equity is declined by every CRRA agent more risk-averse than $\gamma \approx 2.9$, well inside the plausible band; investment-grade bonds only near the top of it.

Appendix B. Computation and replication

The multiplicative index is computed in closed form from Proposition 1. The additive index of the \$100 long position solves $E[\exp\{-\beta \cdot 100(e^{\{X\}} - 1)\}] = 1$ by 100-node Gauss-Hermite quadrature of the log-normal expectation with Brent root-finding, cross-validated by Monte Carlo; the short-position index is recorded non-existent. Rank agreement uses Spearman correlations and an exhaustive discordant-pair scan over all 1,755 within-period pairs. The value/size shape test fits $\alpha_m = a + bd + cd^2$ with period fixed effects and reports the vertex $-b/2c$. The non-parametric estimator of Section 8.4 solves $\hat{w}(\beta) = 1$ on the empirical mgf. All tables and statistics through Section 8 are produced by a single self-contained program from the annual log-return panel; the complete cell-by-cell panel of both indices is reported in Appendix C, which is the panel of record for the summary statistics used throughout Sections 6-8. The monthly pricing tests of Section 9 (Tables 4 and 5) are produced by two further self-contained programs. The first reads the Kenneth French monthly file, constructs the trailing-window predictor from log returns, and runs the Fama-MacBeth, portfolio-sort, and value-forward tests with Newey-West standard errors, writing the per-asset-month predictor panel used in those tests. The second forms the riskiness spread factor and estimates the linear stochastic-discount-factor restriction across the twenty deciles by two-step GMM with a Newey-West weighting matrix, reporting the loading, the implied price of risk, the J -test and the Hansen-Jagannathan distance, with the market excess return run through the identical machinery as a benchmark.

Appendix C. The complete panel

Tables C1 and C2 report every cell of the 27-asset by five-period panel. Table C1 gives the critical CRRA coefficient $\alpha_m = 1 + 2m/s^2$ of Li (2014a); the relative index is recovered as $R_m = 1/\alpha_m$ and is omitted to save space. Table C2 gives the Aumann-Serrano additive index of the corresponding long position, per dollar invested, $R(\varphi - 1)$, obtained by Gauss-Hermite quadrature and Brent root-finding; the additive index of the short position does not exist in any cell (Proposition 2). Assets are grouped as the ten size deciles (smallest to largest), the ten book-to-market deciles (growth to deep value), and seven aggregate classes. Each α_m reads as a CRRA threshold: an agent with relative risk aversion γ holds the asset if and only if $\gamma < \alpha_m$.

Table C1. The complete panel of critical CRRA coefficients $\alpha_m = 1 + 2m/s^2$, by asset and period. Higher α_m means a wider range of CRRA agents will hold the asset; values are large for near-riskless series (Treasury bills especially), consistent with the high- α_m , ordinarily irrelevant cells noted in Section 8.4.

Asset	1928-1945	1946-1965	1966-1985	1986-2005	2006-2025
Size 1	1.90	6.75	3.34	5.78	3.17
Size 2	1.65	7.13	3.88	6.33	4.76
Size 3	1.76	7.75	4.70	8.96	6.26
Size 4	2.03	8.02	4.80	8.27	6.76
Size 5	2.01	7.75	5.42	8.93	5.97
Size 6	2.24	9.49	5.35	11.00	7.81
Size 7	1.88	10.85	5.06	13.58	5.74
Size 8	2.13	10.71	6.92	12.66	6.25
Size 9	1.92	13.41	6.79	13.84	6.13
Size 10	2.25	14.15	6.38	8.88	7.87
Value 1	1.82	10.54	3.56	5.81	6.16
Value 2	3.06	15.85	5.87	11.02	8.27
Value 3	2.23	11.90	5.85	14.35	8.93
Value 4	1.79	14.12	5.12	16.59	7.81
Value 5	1.75	13.05	5.90	20.00	8.38
Value 6	1.68	15.06	9.62	13.15	6.30
Value 7	1.50	8.76	7.97	14.30	2.55
Value 8	1.79	8.98	9.03	12.38	6.33
Value 9	1.95	8.63	8.28	9.67	7.17
Value 10	1.25	5.49	7.58	10.61	2.90
Market	2.14	13.13	6.41	9.96	7.32
S&P 500	3.07	11.20	7.69	10.59	9.44
T-bill 3mo	110.29	376.51	193.38	233.78	86.96
T-bond 10yr	80.42	38.60	15.30	19.85	8.50
Baa bond	19.87	95.22	18.85	52.00	18.68

Asset	1928-1945	1946-1965	1966-1985	1986-2005	2006-2025
Real estate	6.47	19.21	73.89	65.81	15.38
Gold	6.19	3.26	3.32	4.51	7.73

Table C2. The complete panel of the Aumann-Serrano additive index per dollar, $R(\varphi - 1)$, by asset and period. Units are dollars of riskiness per dollar invested in the \$100 long position. Ordinal agreement with Table C1 is near-perfect within each period (Spearman 0.995-0.999; Section 6); the discordant pairs concentrate among the value deciles.

Asset	1928-1945	1946-1965	1966-1985	1986-2005	2006-2025
Size 1	0.5217	0.1407	0.2888	0.1656	0.3102
Size 2	0.6091	0.1338	0.2477	0.1514	0.2041
Size 3	0.5696	0.1228	0.2032	0.1064	0.1542
Size 4	0.4883	0.1182	0.1989	0.1157	0.1430
Size 5	0.4930	0.1227	0.1762	0.1068	0.1617
Size 6	0.4395	0.0999	0.1797	0.0866	0.1231
Size 7	0.5324	0.0874	0.1903	0.0695	0.1682
Size 8	0.4663	0.0886	0.1387	0.0749	0.1539
Size 9	0.5212	0.0706	0.1421	0.0684	0.1571
Size 10	0.4426	0.0670	0.1522	0.1076	0.1215
Value 1	0.5504	0.0901	0.2767	0.1660	0.1544
Value 2	0.3205	0.0599	0.1652	0.0862	0.1154
Value 3	0.4460	0.0803	0.1653	0.0661	0.1068
Value 4	0.5598	0.0675	0.1890	0.0573	0.1236
Value 5	0.5726	0.0719	0.1638	0.0473	0.1150
Value 6	0.5974	0.0625	0.0991	0.0719	0.1531
Value 7	0.6740	0.1089	0.1198	0.0660	0.3881
Value 8	0.5589	0.1043	0.1047	0.0763	0.1526
Value 9	0.5096	0.1094	0.1138	0.0973	0.1339
Value 10	0.8308	0.1728	0.1245	0.0881	0.3393
Market	0.4643	0.0722	0.1506	0.0958	0.1309
S&P 500	0.3180	0.0838	0.1249	0.0893	0.1005
T-bill 3mo	0.0090	0.0026	0.0050	0.0042	0.0114
T-bond 10yr	0.0122	0.0256	0.0634	0.0485	0.1160
Baa bond	0.0488	0.0103	0.0511	0.0183	0.0522
Real estate	0.1537	0.0510	0.0131	0.0148	0.0641
Gold	0.1593	0.3070	0.2868	0.2195	0.1228