

Asset Pricing with Cumulants

Robert Parham

University of Virginia
robertp@virginia.edu

Mihail Velikov

Penn State University
velikov@psu.edu

Han Xiao

CUHK, Shenzhen
xiaohan@cuhk.edu.cn

Abstract

We show that, for any stochastic discount factor (SDF), the log expected gross excess return is an exact sum of joint return-SDF cumulants. The sum nests the familiar covariance channel at order two, co-skewness at order three, and co-kurtosis at order four. When the SDF is affine in a few observable risk factors, one structural parameter per factor prices the entire implied library of cumulant factors. The factor “zoo” is therefore consistent with a low-dimensional SDF. In our tests, three structural parameters generate 31 cumulant prices, and the resulting 28 overidentifying restrictions are not rejected. The restricted model also forecasts held-out returns as well as the unrestricted one. Out of sample on 284 long-short anomaly portfolios, our cumulant models match the Fama–French benchmarks on mean absolute alpha and cut unexplained average returns by a third.

JEL Classification: G11, G12, G14

Keywords: cumulants, co-cumulants, stochastic discount factor, factor zoo, higher moments

1. Introduction

Empirical asset returns depart systematically from the Gaussian benchmark. Option-implied higher moments are large, time-varying, and predictive of future returns and volatility ([Bakshi, Kapadia, and Madan, 2003](#); [Gormsen and Jensen, 2025](#)), and realized higher moments capture a substantial share of the cross-sectional variation in equity returns. [Schneider and Trojani \(2019\)](#) estimate that (almost) model-free higher-moment recovery accounts for a large fraction of priced equity risk, and [Kyle and Todorov \(2023\)](#) document an annualized cumulant risk premium of roughly seven percent on leveraged equity claims. These findings indicate that pricing models built on covariance with the market alone leave systematic structure on the table, concentrated in the higher-order joint behavior of returns and the stochastic discount factor.

A natural question is how to organize this higher-order structure in a way that respects the Euler equation, and that nests the familiar covariance channel at the lowest order. The existing higher-moment cross-sectional literature has approached this question, for the most part, one moment at a time. [Kraus and Litzenberger \(1976\)](#) append co-skewness to the CAPM, [Harvey and Siddique \(2000\)](#) estimate conditional co-skewness premia, [Dittmar \(2002\)](#) adds co-kurtosis through a polynomial pricing kernel in the spirit of the orthogonal-polynomial kernel approximations of [Chapman \(1997\)](#), and [Chabi-Yo \(2012\)](#) introduces stochastic skewness. Each of these contributions delivers one new moment with its own kernel specification, estimation procedure, and test design.

In this paper, we argue that *cumulants*, rather than raw moments, are the natural object for this purpose. First, joint cumulants vanish when any subset of arguments is independent of the rest, so only genuine dependence between returns and the SDF contributes to pricing. Second, cumulants are additive for independent random variables. And finally, the cumulant generating function converts the multiplicative structure of the Euler equation into a linear one in log space.

Beyond organizing higher-order co-movement between asset prices and the SDF, the same cumulant structure speaks to a second, much-debated issue: the dimensionality of the cross section of stock returns. Hundreds of long-short signals have been documented in the empirical asset pricing literature (Cochrane, 2011; Harvey, Liu, and Zhu, 2016), and a basic question is how many independent priced factors actually underlie them. Casting pricing in cumulant space lets us approach this structurally rather than statistically, where instead of selecting factors empirically from a large additive specification, we ask what the Euler equation itself implies about the dimensionality of priced risk.

Our starting point is the standard no-arbitrage condition $\mathbb{E}[MR_i] = 1$. We first rewrite it as $\log \mathbb{E}[e^{r_i - m}] = 0$, with $r_i \equiv \log R_i$, $m \equiv -\log M = \log M^{-1}$, and we call the log-inverse-pricing kernel m the SDF throughout. We hence work with the cumulant generating function of excess log returns evaluated at one, the $\mathbb{K}$ -expectation, $\mathbb{K}[X] \equiv \log \mathbb{E}[e^X]$. The standard no-arbitrage condition is rewritten as $\mathbb{K}[r_i - m] = 0$. Applied to the excess log return, $\mathbb{K}[r_i - r_f]$ equals the log of the expected gross excess return; we call it the $\mathbb{K}$ -expected excess log return. This is not the arithmetic mean of $r_i - r_f$, but instead the object whose pricing is equivalent to the standard Euler equation, and the natural pivot from the Euler equation in gross return space to a structured representation in the cumulant log-space.

The exact log-space linear identity at the heart of the framework states that

$$\mathbb{K}[r_i - r_f] = \sum_{a,b \geq 1} \frac{(-1)^{b+1}}{a! b!} \kappa_{a,b}(r_i, m), \quad (1)$$

where $\kappa_{a,b}(r_i, m)$ is the joint cumulant of order a in r_i and order b in the SDF m , and the coefficients are known closed-form combinatorial constants implied by the cumulant expansion rather than free parameters. The identity holds for any admissible SDF satisfying the Euler equation and any return distribution for which the joint moment generating function of (r_i, m) is finite at the unit argument. The equivalence is established as Theorem 1 in Section 3.

At truncation order N , the pricing equation (1) lives in a finite-dimensional co-cumulant

space whose dimension grows only quadratically in the order, with coefficients that are given by the theory in closed form rather than estimated. A small number of priced factors per order, with theory-dictated weights, replaces the open-ended menu of free factor loadings in additive linear specifications.

Furthermore, within each total order $n = a + b$, the individual mixed co-cumulants in Equation (1) can be aggregated into a single composite beta and a single price of risk. The composite beta collapses the within-order co-cumulant loadings of asset i into one scalar exposure to the order- n cumulant of the SDF. At order two, it reduces to the standard covariance beta; order three contributes one composite co-skewness loading, order four one composite co-kurtosis loading, and so on. The pricing equation, therefore, has one beta-price-of-risk term per order, which is a small number of sufficient statistics compared to the menu of individual co-cumulants, and the basis for the structural compression of the factor zoo that we exploit empirically below.

Mapping the log-space identity back to gross return space yields a multiplicative factor structure for gross expected returns. Exponentiating (1) and rearranging shows that gross expected returns are an exact product of order-level factors, one per cumulant order, rather than a sum. This is the multiplicative cross-sectional counterpart to the additive structure of [Ross \(1976\)](#) arbitrage pricing theory and of standard linear factor models. We argue that the familiar additive representation in simple returns can be interpreted as a low-order log-linearization of this Euler-implied multiplicative structure, derived in Section 3.

The theory so far conditions on the latent SDF. Three further propositions close the gap to estimation. If the SDF is affine in a small set of observable risk factors (e.g., affine in the market return in the leading single-factor case, or affine in a set of factors as in, say, [Kozak and Nagel \(2023\)](#), and most of the asset pricing literature), then the multilinearity property of joint cumulants decomposes every mixed co-cumulant exactly, with no approximation, into joint cumulants of the return with the risk factors (Proposition 2). We call these joint-

cumulant terms of returns with the risk factors the *cumulant factors*.¹ The affine structure simultaneously pins the prices of the entire expansion down to simple products of the SDF loadings. The price coefficients of the whole mixed-co-cumulant library, to any expansion order, are determined by only one parameter per observable risk factor (Proposition 3).

For the six Fama–French risk factors at cumulant order four, six parameters hence determine the structural prices of 116 cumulant factors, a heavily overidentified structure. These exposures can be expressed in estimable beta coordinates that coincide with the familiar regression betas at order two and extend to all higher orders. The overidentification is the framework’s central testable prediction, and we take it to the data directly: at a three-factor span and fourth order, three parameters determine 31 cumulant prices, leaving 28 overidentifying restrictions that a bootstrapped minimum-distance test can evaluate without asymptotic approximations. In our reduced-form pricing tests, by contrast, we do not impose the restriction: the estimated specification prices a small semi-aggregated subset of exposures (eight at six factors and order four) with free Fama–MacBeth prices.

We take the identity to data to learn which cumulant risks are priced across the zoo of anomalies. Our test assets are the 284 collateralized long-short anomaly portfolios of [Chen and Zimmermann \(2022\)](#), and we use a small set of risk factors: the log excess market return and collateralized log versions of the Fama–French spread factors. For each asset, we estimate its joint cumulants with these factors from exponentially weighted moments, compress them order by order, and price the resulting order-level exposures in monthly Fama–MacBeth regressions in log space ([Fama and MacBeth, 1973](#)). The tests are out of sample in the pricing dimension over 1978 to 2024: every ingredient of the month $t + 1$ forecast uses only information through month t . We compare against the standard benchmarks: the rolling-beta CAPM ([Sharpe, 1964](#); [Lintner, 1965](#); [Black, 1972](#)), the [Fama and French \(1993\)](#) three-factor model, and the Fama–French six-factor model, each estimated in the usual way and compared in the same metric space.

¹ Reserving *priced* for the ones that earn a premium in the cross section.

The reorganization delivers its gains at no cost to pricing accuracy. On mean absolute alpha, the models are statistically tied: on the common-support panel that serves as our default baseline, where all models are evaluated on one identical asset-month panel, the richest cumulant model edges raw FF6 descriptively (13.09 versus 13.40 basis points per month), but a stationary block bootstrap that respects month-level dependence and cross-sectional error correlation finds the pairwise differences indistinguishable from zero in either direction. The gains appear instead in the tails and in significance. Alpha significance improves monotonically with cumulant order: the mean-squared alpha t -statistic falls from 1.47 under raw FF6 to 1.26 under the fourth-order six-factor model, and the share of assets with $|t_\alpha| > 1.96$ falls from 10.6% to 9.5%. The mean Fama–MacBeth intercept, the average return the cross-sectional regressions cannot attribute to priced exposures, drops from 16.9 to 11.1 basis points, a one-third reduction in unexplained average returns. These improvements are distributional, spread broadly across the cross section rather than reflecting a point ranking on a single error statistic.

The paper’s central test then asks whether the structure behind these models is supported by the data. The unrestricted model estimates 31 free prices of risk each month; the affine null says the same 31 prices are generated by three structural loadings. We estimate the 31 prices freely, month by month, on a quintile-portfolio panel of roughly 1,400 assets, fit the three-parameter restricted model to the estimated price vector by weighted least squares (a minimum-distance estimation), and calibrate the resulting price residual with circular block bootstraps, including a null-imposed design in which pseudo-panels satisfy the restriction by construction and the entire pipeline is rerun on each draw. We report three findings. The restrictions are not rejected in price space: the observed aggregate price residual is smaller than what the null-imposed pipeline generates from estimation noise alone. The in-sample fit gain from freeing all 31 prices is almost exactly what the same relaxation extracts from panels where the restriction holds (overfitting, not signal). And on held-out data, the restricted three-parameter model forecasts as well as or better

than the 31-parameter unrestricted model, decisively so for the (MKT, RMW, CMA) span. The data, in short, cannot distinguish the large cumulant-factor library from the handful of structural parameters the theory says generate it.

We then run a decomposition to locate where the structure does its work. We group the 284 signals into seventeen anomaly families and ask, signal by signal, which anomalies the rich cumulant model still fails to price. Eighteen of the 284 signals retain a significant $\mathbb{K}$ -Beta[6,4] alpha, most of them momentum- and seasonality-type persistence signals; raw FF6 fails on fifteen of the same eighteen, so the survivors appear to reflect risk factors outside the FF6 span rather than an artifact of the cumulant machinery. We further present which anomaly families load on which risk factor and cumulant exposure, and find that the expected families load on the expected FF6 factors, and the 3rd and 4th order cumulants are highly active for many of the anomaly families.

We make three theoretical contributions. First, we prove that the log-space identity in Equation (1) is an exact Euler-equation equivalence with explicit closed-form coefficients, formalizing and cross-sectionally extending the coentropy concept introduced by [Backus, Boyarchenko, and Chernov \(2018\)](#) for the term structure of asset prices. The identity reorganizes higher-order pricing in that instead of appending one moment at a time, every co-cumulant channel enters a single equation with a weight the Euler equation pins down in closed form.

Second, we carry the theory to observables. When the SDF is affine in a small set of risk factors, we derive an exact multilinear decomposition onto the observable risk factors, the structural restriction that ties the prices of the whole cumulant-factor library to one parameter per factor, and the estimable $\mathbb{K}$ -beta coordinates. The restriction turns the factor zoo from a list of free prices into a few structural parameters and generates the overidentifying restrictions that we then take to data.

Third, we derive two representation results from the identity. Within each total order, the mixed co-cumulants of asset returns and the SDF collapse into one composite beta and

one price of risk, a one-beta-per-order representation that nests the covariance channel at order two and connects to the co-skewness work of [Kraus and Litzenberger \(1976\)](#); [Harvey and Siddique \(2000\)](#) and the co-kurtosis work of [Dittmar \(2002\)](#); [Chabi-Yo \(2012\)](#), which appear as truncations at the corresponding orders. And in gross return space, the implied factor structure for expected returns is multiplicative rather than additive, the counterpart to [Ross \(1976\)](#) and standard linear factor models, and distinct from the time-series operator decomposition of [Hansen and Scheinkman \(2009\)](#).

Our empirical contributions parallel the theory. First, we price the $\mathbb{K}$ -beta exposures of 284 long-short anomaly portfolios out of sample and find a statistical tie with the Fama–French six-factor benchmark on mean absolute alpha, alongside consistent gains in alpha significance, in the tails of the alpha distribution, and in the Fama–MacBeth intercept. Second, we test the framework’s 28 overidentifying restrictions directly and cannot reject them, while the restricted model matches or beats the unrestricted one on held-out data. Third, a signal-level decomposition localizes the residual factor zoo to a short list of momentum-type signals.

The rest of the paper is organized as follows. Section 2 relates the framework to the literature. Section 3 derives the cumulant-space pricing equation for a general SDF: the exact identity, the order-level composite-beta representation, and the multiplicative factor structure. Section 4 specializes to the affine case, derives the structural price restriction and the estimable $\mathbb{K}$ -beta coordinates, and states the overidentifying restrictions that the empirical work tests. Section 5 presents the empirics: the out-of-sample pricing tests against the CAPM, FF3, and FF6 benchmarks, the test of the overidentifying restrictions, and the signal-level decomposition. Section 6 concludes. Proofs are collected in Appendix A. An Internet Appendix collects supporting theory (Section IA1), implementation details (Section IA2), alternative SDF constructions (Section IA3), additional diagnostics (Section IA4), a notation glossary (Section IA5), and the signal-to-family mapping (Section IA6).

2. Related Literature

Our framework connects to several strands of the literature. Equation (1) formalizes and cross-sectionally extends the coentropy concept introduced by [Backus et al. \(2018\)](#) for the term structure of asset prices and further developed by [Chabi-Yo and Colacito \(2019\)](#) in the international setting. Coentropy is the bivariate joint-cumulant object that those papers use to study term-structure pricing; our identity is the corresponding Euler-equation equivalence theorem, stated with explicit existence conditions, applied to the cross section of expected returns. Two related strands use the cumulant generating function of the SDF itself: [Backus, Chernov, and Zin \(2014\)](#) decompose the entropy of the marginal of m across representative-agent models, and [Backus, Chernov, and Martin \(2011\)](#) use the joint (r, m) structure to recover disaster probabilities from index options.

The framework also sits within a broader tradition of entropy- and cumulant-based methods for stochastic discount factors. [Backus et al. \(2014\)](#) use the entropy of m , the marginal cumulant generating function, to discriminate among representative-agent models; [Alvarez and Jermann \(2005\)](#) bound the persistent component of marginal utility via SDF entropy; [Bekaert and Engstrom \(2017\)](#) use entropy diagnostics in a habits-with-bad-environment-good-environment economy; and [Ghosh, Julliard, and Taylor \(2017\)](#) develop information-theoretic SDF bounds that nest the [Hansen and Jagannathan \(1991\)](#) variance bounds. The exponential-affine SDF of [Bansal and Yaron \(2004\)](#) long-run-risks specification is a leading parametric instance of this class, and [Martin \(2017\)](#) uses the SVIX bound to deliver a model-free lower bound on the equity premium.

Our contribution is adjacent in spirit but distinct in focus: we work on the joint cumulants of (r_i, m) rather than the marginal of m alone, and we deliver an equality for cross-sectional pricing rather than a bound on aggregate quantities. Operating in the cumulant generating function is also robust to the concern raised by [Martin and Shi \(2026\)](#) that moments of the SDF may diverge before reaching second order: the joint moment generating

function of (r_i, m) can remain finite at the unit argument even when individual moments of arbitrary order do not, so cumulant-based pricing avoids the fragility of Taylor-expansion-in-moments approaches. A separate strand decomposes the SDF itself multiplicatively across maturities via spectral methods ([Hansen and Scheinkman, 2009](#); [Hansen, 2012](#)); our multiplicative gross-return structure in Equation (13) is the cross-sectional analogue at a fixed horizon, distinct from those time-series operator decompositions.

The composite-beta hierarchy provides a natural home for the higher-moment cross-sectional literature discussed above. [Kraus and Litzenberger \(1976\)](#) introduce a three-moment CAPM with co-skewness priced alongside covariance; [Harvey and Siddique \(2000\)](#) show that conditional co-skewness helps explain the cross section of expected returns; [Dittmar \(2002\)](#) extends the pricing kernel to include kurtosis preference and finds that a nonlinear four-moment SDF improves cross-sectional fit; [Chabi-Yo \(2012\)](#) prices stochastic skewness; [Bakshi et al. \(2003\)](#) deliver model-free skewness and kurtosis from index options; [Schneider and Trojani \(2019\)](#) show that higher-moment recovery accounts for a substantial share of return variation; and [Gormsen and Jensen \(2025\)](#) document substantial time variation in higher-moment risk and its asset-pricing implications. Each of these contributions introduces one moment at a time, with its own specification and test design. Our decomposition derives the full hierarchy from a single Euler-equation identity and nests each specification as a special case at the corresponding truncation order. The order-by-order structure also provides a natural stopping rule: one can test directly whether adding the next order materially reduces pricing errors.

Four CGF-based works deserve careful distinction. [Martin \(2013\)](#) expands consumption-based asset pricing in higher cumulants of the *aggregate* consumption growth process under Epstein–Zin preferences; that representation is univariate and aggregate, while ours is bivariate in (r_i, m) and applies to any SDF and any assets. [Snow \(1991\)](#) provides an early diagnostic showing that cumulants of the SDF appear in the Euler equation under a Taylor expansion in moments; our identity is exact and avoids the moment-existence

problems that motivate [Martin and Shi \(2026\)](#). [Orłowski, Sali, and Trojani \(2019\)](#) use the joint CGF of (r, m) to derive bounds on the dispersion of admissible pricing kernels via Jensen’s inequality; ours is an equality, not a bound. Most recently, [Kyle and Todorov \(2023\)](#) construct a cumulant risk premium on leveraged equity claims, where the relevant exposures take the form β^n (leverage raised to the n -th power on a single underlying) rather than our composite $B_i^{(n)}$, which aggregates mixed co-cumulant betas across distinct (a, b) pairs at each total order. We view their findings as direct evidence that the higher-order cumulant channels are economically important; our framework provides the structural decomposition that organizes such evidence across general assets.

The decomposition also speaks to the dimensionality question. Standard practice runs benchmark horse races against the CAPM, the [Fama and French \(1993\)](#) three-factor model, [Carhart \(1997\)](#), the [Fama and French \(2015\)](#) five-factor model, and the [Hou, Xue, and Zhang \(2015\)](#) q -factor model; recent work tackles dimensionality directly through PCA, penalized regressions, machine learning, and Bayesian model averaging ([Feng, Giglio, and Xiu, 2020](#); [Kozak, Nagel, and Santosh, 2020](#); [Kelly, Pruitt, and Su, 2019](#); [Gu, Kelly, and Xiu, 2020](#); [Bryzgalova, Huang, and Julliard, 2023](#); [Kozak and Nagel, 2023](#); [Borri, Chetverikov, Liu, and Tsyvinski, 2026](#)). We provide a complementary structural answer: under the affinity condition this literature works with, the objects carrying structural prices are not the spanning factors themselves but the larger library of cumulant factors of Proposition 3, with weights pinned down in closed form and prices tied together by only one parameter per spanning factor. Dimensionality, therefore, splits into span and order, and a small SDF span is compatible with a combinatorially rich set of cumulant factors. In that sense, the decomposition provides a structural compression map from the factor zoo to a smaller set of underlying parameters, complementary to, not competing with, the statistical compression results in the literature; we develop this reading against the specific findings of that literature in Section 4.5.

The starting point, the Euler equation and a positive SDF, is the same as in the no-

arbitrage and consumption-based asset pricing tradition (Lucas, 1978; Breeden, 1979; Hansen and Richard, 1987; Hansen and Jagannathan, 1991; Cochrane, 2005). We do not propose a new pricing kernel. Any SDF that satisfies the Euler equation, including consumption-based, intermediary (He and Krishnamurthy, 2013; Adrian, Etula, and Muir, 2014), and reduced-form specifications, implies the co-cumulant pricing decomposition described above.

3. Theory: The General Case

We work in unconditional notation throughout this section. When we take the objects to the data, we interpret them conditionally and estimate them in rolling windows using information available through $t - 1$ only. We denote an asset's gross return by $R_i > 0$, the gross risk-free return by $R_f > 0$, and the pricing kernel by $M > 0$, with log counterparts $r_i \equiv \log R_i$, $r_f \equiv \log R_f$, and $m \equiv -\log M$ with m denoted the SDF throughout.

The law of one price implies the existence of a positive pricing kernel satisfying

$$\mathbb{E}[MR_i] = 1, \quad (2)$$

for every traded payoff and, for the risk-free asset,

$$\mathbb{E}[MR_f] = 1. \quad (3)$$

This is the standard starting point of the no-arbitrage and SDF literature (Lucas, 1978; Breeden, 1979; Hansen and Richard, 1987; Hansen and Jagannathan, 1991; Cochrane, 2005). It nests the consumption-based case, in which M is the discounted ratio of marginal utilities, and does not commit us to a single structural kernel. Our question is what these Euler restrictions imply when expressed in log-cumulant language.

Because $M = e^{-m}$ and $R_i = e^{r_i}$, the Euler restrictions are equivalent to

$$\mathbb{K}[r_i - m] = 0, \quad \mathbb{K}[r_f - m] = 0, \quad (4)$$

where, recall, $\mathbb{K}[X] \equiv \log \mathbb{E}[e^X]$ is the $\mathbb{K}$ -expectation, the cumulant generating function evaluated at one. We develop the implications of Equation (4) in the rest of this section.

3.1. Euler-equation equivalences

For any scalar random variable X , define its *cumulant generating function*

$$K_X(s) \equiv \log \mathbb{E}[e^{sX}] = \sum_{n=1}^{\infty} \frac{s^n}{n!} \kappa_n(X),$$

and for a pair (X, Y) define its joint cumulant generating function

$$K_{X,Y}(s_1, s_2) \equiv \log \mathbb{E}[e^{s_1 X + s_2 Y}], \quad \kappa_{a,b}(X, Y) \equiv \left. \frac{\partial^{a+b}}{\partial s_1^a \partial s_2^b} K_{X,Y}(s_1, s_2) \right|_{(0,0)}.$$

The first index in $\kappa_{a,b}$ always counts copies of the asset return, and the second counts copies of the SDF.

Theorem 1 (Euler-equation equivalence theorem). *Assume the joint moment generating function of (r_i, m) is finite at the relevant arguments, so that the cumulant expansions below converge. For any traded asset i , the following statements are equivalent:*

1. the Euler equation $\mathbb{E}[MR_i] = 1$;
2. the log-represented Euler equation $\log \mathbb{E}[e^{r_i - m}] = 0$;
3. the $\mathbb{K}$ -expectation equation $\mathbb{K}[r_i - m] = 0$;
4. the cumulant expansion $\sum_{n \geq 1} \kappa_n(r_i - m)/n! = 0$;
5. the co-cumulant expansion $\sum_{a,b \geq 0} (-1)^b \kappa_{a,b}(r_i, m)/(a! b!) = 0$;
6. the mixed co-cumulant pricing identity

$$\mathbb{K}[r_i - r_f] = \sum_{a,b \geq 1} c_{a,b} \kappa_{a,b}(r_i, m), \quad c_{a,b} \equiv \frac{(-1)^{b+1}}{a! b!}; \quad (5)$$

7. the beta-lambda representation

$$\mathbb{K}[r_i - r_f] = \sum_{a,b \geq 1} c_{a,b} \beta_i^{(a,b)} \lambda_{a+b}, \quad \beta_i^{(a,b)} \equiv \frac{\kappa_{a,b}(r_i, m)}{\kappa_{a+b}(m)}, \quad \lambda_n \equiv \kappa_n(m); \quad (6)$$

8. the gross-space representation

$$\mathbb{E}[R_i] = R_f \exp\left(\sum_{n \geq 2} B_i^{(n)} \lambda_n\right) = R_f \prod_{n \geq 2} \Lambda_n^{B_i^{(n)}}, \quad \Lambda_n \equiv e^{\lambda_n}, \quad (7)$$

where $B_i^{(n)}$ denotes the order-level composite beta defined below.

We collect the proof in Appendix A.

3.2. The core pricing equation in mixed co-cumulants

Theorem 1 part (6) implies that all priced risk is spanned by mixed co-cumulants of the asset return and the SDF. Once we impose the risk-free restriction, pure-return cumulants and pure-SDF cumulants cancel. What remains are the dependence terms that tell us how the payoff loads on discounting states.

For any truncation order $N \geq 2$, we define

$$\zeta_i^{(N)} \equiv \sum_{\substack{a+b \geq N+1 \\ a, b \geq 1}} c_{a,b} \kappa_{a,b}(r_i, m). \quad (8)$$

Then

$$\mathbb{K}[r_i - r_f] = \sum_{\substack{a+b \leq N \\ a, b \geq 1}} c_{a,b} \kappa_{a,b}(r_i, m) + \zeta_i^{(N)}. \quad (9)$$

Up to the truncation residual $\zeta_i^{(N)}$, pricing therefore lives in a $D_N = \binom{N}{2}$ -dimensional mixed co-cumulant space. We hypothesize that a few low orders suffice, as we discuss below. Through the fourth order, the mixed co-cumulants admit a simple interpretation. Table 1 reports the six pricing factors that span the fourth-order model.

Collecting the terms order by order, the fourth-order truncation is

$$\mathbb{K}[r_i - r_f]^{(4)} = \underbrace{\kappa_{1,1}(r_i, m)}_{\text{order 2}} + \underbrace{\frac{1}{2}\kappa_{2,1}(r_i, m) - \frac{1}{2}\kappa_{1,2}(r_i, m)}_{\text{order 3}} + \underbrace{\frac{1}{6}\kappa_{3,1}(r_i, m) - \frac{1}{4}\kappa_{2,2}(r_i, m) + \frac{1}{6}\kappa_{1,3}(r_i, m)}_{\text{order 4}}, \quad (10)$$

where truncating at order two or three simply drops the order-four (or order-three and four) brace. The covariance term is the order-two benchmark. The higher-order terms are

Table 1. Mixed co-cumulants through fourth order

Basis	Mixed co-cumulant	Order	Coefficient	Interpretation
$e_{1,1}$	$\kappa_{1,1}(r_i, m)$	2	+1	Return level with SDF level.
$e_{2,1}$	$\kappa_{2,1}(r_i, m)$	3	$+\frac{1}{2}$	Return variance with SDF level.
$e_{1,2}$	$\kappa_{1,2}(r_i, m)$	3	$-\frac{1}{2}$	Return level with SDF variance.
$e_{3,1}$	$\kappa_{3,1}(r_i, m)$	4	$+\frac{1}{6}$	Return skewness with SDF level.
$e_{2,2}$	$\kappa_{2,2}(r_i, m)$	4	$-\frac{1}{6}$	Return variance with SDF variance.
$e_{1,3}$	$\kappa_{1,3}(r_i, m)$	4	$+\frac{1}{6}$	Return level with SDF skewness.

the remaining factors implied by the same Euler equation once we leave Gaussianity. Each entry in Table 1 has a transparent covariance-like reading. Order two, $\kappa_{1,1}(r_i, m)$, is just the covariance of the return with the SDF. The third-order terms $\kappa_{2,1}$ and $\kappa_{1,2}$ are covariances built from squared innovations (the return variance's covariance with the SDF level and the return level's covariance with the SDF variance), and the fourth-order terms layer in cubed innovations the same way. In the $\mathbb{K}$ -beta coordinates of Section 4.3, these higher co-cumulants become exact covariances of the corresponding power transforms, which is what makes them estimable as ordinary multivariate betas.

3.3. Order-level pseudo-prices and composite betas

The co-cumulant representation is exact, but empirical work benefits from a normalization that packages the within-order detail. For each (a, b) with $a, b \geq 1$, define the normalized co-cumulant beta

$$\beta_i^{(a,b)} \equiv \frac{\kappa_{a,b}(r_i, m)}{\kappa_{a+b}(m)}.$$

The corresponding term in the pricing equation can then be written as

$$c_{a,b} \kappa_{a,b}(r_i, m) = \beta_i^{(a,b)} \pi_{a,b}, \quad \pi_{a,b} \equiv c_{a,b} \kappa_{a+b}(m).$$

The object $\pi_{a,b}$ is a *pseudo-price of co-cumulant risk*: it is the theory-implied scaling attached to the normalized exposure, but it is not a dollar price and should not be interpreted as one.

Within a fixed total order $n = a + b$, all pseudo-prices share the same SDF cumulant $\lambda_n \equiv \kappa_n(m)$. The only differences across the (a, b) terms are the combinatorial coefficients $c_{a,b}$. This permits an exact within-order collapse. We define the composite beta

$$B_i^{(n)} \equiv \sum_{\substack{a+b=n \\ a,b \geq 1}} c_{a,b} \beta_i^{(a,b)} = \sum_{\substack{a+b=n \\ a,b \geq 1}} c_{a,b} \frac{\kappa_{a,b}(r_i, m)}{\kappa_n(m)}. \quad (11)$$

The pricing identity becomes

$$\mathbb{K}[r_i - r_f] = \sum_{n \geq 2} B_i^{(n)} \lambda_n. \quad (12)$$

At order four, the pricing rule collapses to

$$\mathbb{K}[r_i - r_f]^{(4)} = B_i^{(2)} \lambda_2 + B_i^{(3)} \lambda_3 + B_i^{(4)} \lambda_4 + \zeta_i^{(4)}.$$

This is the low-dimensional order-level factor model we take to the data. The truncation residual $\zeta_i^{(4)} = \sum_{n \geq 5} B_i^{(n)} \lambda_n$ collects orders five and above. It is small for a concrete reason: monthly $r_i - r_f$ and m are small in magnitude, so the order- n joint cumulants shrink quickly in n , and the combinatorial weights $c_{a,b}$ carry factorial denominators that shrink faster still. The product of the two makes the order-five and order-six contributions empirically negligible; untabulated order-five and order-six diagnostics confirm this directly rather than assuming it.

The collapse from the full (a, b) system to the order-level system has an important economic implication. The relevant question is whether there are economically important order-level pricing factors once we impose the Euler-equation weights, not how many named higher-moment statistics can be written down.

Remark (Nesting the higher-moment CAPM). At $D = 1$ with the single factor (the market), the composite betas are the familiar higher-moment objects: $B_i^{(2)}$ is the market beta of the CAPM, $B_i^{(3)}$ is, up to scale, the coskewness loading of [Kraus and Litzenberger \(1976\)](#) and [Harvey and Siddique \(2000\)](#), and $B_i^{(4)}$ the cokurtosis loading of [Dittmar \(2002\)](#). Truncating Equation (12) at order P reproduces those models as exact special cases; what the Euler

equation adds is the relative locked pricing weights $c_{a,b}$ across orders, which those papers leave free.

3.4. Back to gross-space: A multiplicative factor structure

The order-level pricing rule is native to log space. Exponentiating it returns the model to gross-space. Since

$$e^{\mathbb{E}[r_i - r_f]} = \mathbb{E}\left[\frac{R_i}{R_f}\right],$$

the exact gross-space counterpart of Equation (12) is

$$\mathbb{E}[R_i] = R_f \exp\left(\sum_{n \geq 2} B_i^{(n)} \lambda_n\right) = R_f \prod_{n \geq 2} \Lambda_n^{B_i^{(n)}}, \quad \Lambda_n \equiv e^{\lambda_n}. \quad (13)$$

This is a multiplicative factor structure for gross expected returns, and it recasts how we read the factor zoo. Many different portfolio combinations can provide noisy or redundant additive approximations to the same underlying multiplicative structure in gross-space, while the cumulant expansion in log space makes the primitive pricing factors explicit. The contrast is with the usual additive expected-return language of CAPM and multifactor linear models (Sharpe, 1964; Lintner, 1965; Black, 1972; Fama and French, 1993, 2015).

3.5. Implications and stopping rule

The cumulant expansion yields three immediate implications. First, all priced risk is spanned by mixed co-cumulants. Second, the spanned risk is low-dimensional once the model is truncated at a modest order. Third, the gross-space representation implied by the law of one price is multiplicative, not additive.

The theory alone does not decide where to truncate. That is an empirical question. We take fourth order as our main stopping point: it retains the first six mixed factors, preserves a compact order-level representation, and leaves orders five and six available as direct magnitude checks in the data. Orders two through four are our main estimated models; orders five and six serve only as diagnostics that justify the stopping rule.

4. Theory: The Affine Case

4.1. The exact affine decomposition

The results of Section 3 condition on the latent SDF: the cumulant factors $\kappa_{a,b}(r_i, m)$ and the prices $\lambda_n = \kappa_n(m)$ are functions of a variable no one observes. The three results of this section close the gap between Theorem 1 and anything estimable. We replace the latent m with a small vector of observable risk factors and establish three results. First, if the SDF is affine in the risk factors, every mixed co-cumulant in the pricing identity decomposes *exactly* into joint cumulants of the return with the risk factors (Proposition 2). Second, the affine structure pins the prices of the entire (combinatorially much larger) library of cumulant terms down to monomials in the D SDF loadings, products of powers of the loadings, $\eta_{p,q} = \prod_d b_d^{q_d}$ (Proposition 3). Third, an estimable normalization of the observable exposures, the $\mathbb{K}$ -beta transform, generalizes the multivariate regression beta to higher orders and turns the structural prices into monomials in the factor-level covariance prices (Proposition 4). Together, these results are the formal bridge between the theorem and the reduced-form implementation of Section 5.

Let $f = (f_1, \dots, f_D)'$ be a vector of observable risk factors in units compatible with log returns; Section 5.1.2 constructs them as the log collateralized Fama–French factor payoffs. We fix the terminology used in the rest of the paper precisely. The D components of f are the *risk factors*, the variables that span the SDF; they are few. The joint-cumulant terms of returns with these factors that the expansion below generates are the *cumulant factors* $Q_{i,p,q}$; they are combinatorially many. A cumulant factor is *priced* when it commands a nonzero premium in the cross section. Risk factors span the SDF; cumulant factors are the objects that carry the cross-sectional prices.

Assumption 1. *The SDF is affine in the risk factors,*

$$m = a + \sum_{d=1}^D b_d f_d = a + b' f, \quad \text{for } a \in \mathbb{R}, \quad b = (b_1, \dots, b_D)' \in \mathbb{R}^D, \quad (14)$$

and the joint cumulants of (r_i, f) of all relevant orders exist.

For $D = 1$ with f_1 being the log excess market return, Equation (14) is the affine-in-market specification of the introduction. In the conditional implementation, the coefficients are $\mathcal{F}_t$ -measurable, $m_{t+1} = a_t + b'_t f_{t+1}$, and every statement below holds conditionally with t subscripts. The assumption restricts the projection structure of the SDF, not the distribution of returns: nothing below invokes Gaussianity or any other distributional restriction.

We index cumulant terms as in the empirical sections: $p \geq 1$ counts copies of the return, and the vector $q = (q_1, \dots, q_D)$ of nonnegative integers records the factor copies (q_d counts copies of factor d), with $|q| \equiv \sum_d q_d$ and $b^q \equiv \prod_d b_d^{q_d}$. Define

$$\kappa_{p,q}(r_i, f) \equiv \kappa(\underbrace{r_i, \dots, r_i}_{p \text{ copies}}, \underbrace{f_1, \dots, f_1}_{q_1 \text{ copies}}, \dots, \underbrace{f_D, \dots, f_D}_{q_D \text{ copies}}), \quad c_{p,q} \equiv \frac{(-1)^{|q|+1}}{p! \prod_d q_d!}.$$

For $D = 1$, the pair (p, q) is the (a, b) of Theorem 1 and $c_{p,q}$ reduces to $c_{a,b}$.

Proposition 2 (Multilinear cumulant decomposition). *Under Assumption 1, for every $p \geq 1$ and $h \geq 1$,*

$$\kappa_{p,h}(r_i, m) = \sum_{|q|=h} \frac{h!}{\prod_d q_d!} b^q \kappa_{p,q}(r_i, f), \quad (15)$$

and the pricing identity, Equation (5), becomes

$$\mathbb{E}[r_i - r_f] = \sum_{p \geq 1, |q| \geq 1} \eta_{p,q} Q_{i,p,q}, \quad Q_{i,p,q} \equiv c_{p,q} \kappa_{p,q}(r_i, f), \quad \eta_{p,q} \equiv b^q. \quad (16)$$

The SDF intercept a does not enter.

The decomposition is *exact*, not an approximation: joint cumulants are multilinear in each argument, so substituting Equation (14) into the h SDF arguments of $\kappa_{p,h}(r_i, m)$ and collecting symmetric terms yields Equation (15) with no remainder. This is the multilinearity bridge invoked at the introduction: under the affine assumption, the observable exposures $Q_{i,p,q}$ are exactly the building blocks of the theorem's expansion, and the latent-SDF and observable-factor representations are the same pricing equation written in two

coordinate systems. Truncation carries over verbatim: each substitution preserves total order, so an order- $(p + h)$ SDF term maps only into order- $(p + |q|) = (p + h)$ factor terms, and truncating Equation (16) at total order N leaves exactly the residual $\zeta_i^{(N)}$ of Equation (8).

4.2. The structural restriction on cumulant prices

Read as a cross-sectional pricing equation, Equation (16) is linear: the exposures $Q_{i,p,q}$ are asset-level regressors and the $\eta_{p,q}$ are prices. In reduced form, the prices may be estimated freely. The affine structure says they are anything but free.

Proposition 3 (Structural price restriction). *Under Assumption 1, the pricing equation, Equation (16), holds with*

$$\eta_{p,q} = b^q = \prod_{d=1}^D b_d^{q_d} \quad \text{for every } p \geq 1, |q| \geq 1. \quad (17)$$

Consequently:

1. Prices depend on (p, q) only through q : all terms sharing a factor-copy pattern q carry one common price, regardless of the return order p ;
2. The term library $\mathcal{B}(D, P) = \{(p, q) : p \geq 1, |q| \geq 1, p + |q| \leq P\}$ contains

$$J(D, P) = \sum_{h=1}^{P-1} (P - h) \binom{D + h - 1}{h} \quad (18)$$

terms, all of whose prices are determined by the D parameters $b_1, \dots, b_D$.

At order two, the restriction is invisible: D premia for D betas, which is the just-identified textbook case. The content begins at order three, where the premia stop being free. In the conditional implementation, the restriction reads $\eta_{p,q,t} = \prod_d b_{d,t}^{q_d}$. The compression in part (2) is severe. For $D = 1$, $J(1, P) = \binom{P}{2}$, recovering the dimension D_N of Section 3.2, and Equation (17) collapses to powers of a single loading, $\eta_{p,q} = b^q$, the cross-asset generalization of the leverage-power exposures of Kyle and Todorov (2023). For the six Fama–French factors at order four, the structural price coefficients of the $J(6, 4) = 116$

cumulant factors are determined by six parameters, leaving 110 overidentifying restrictions.

A cumulant factor carries a nonzero *structural price coefficient*, $\eta_{p,q} = \prod_d b_d^{q_d}$, under the affine null whenever all $b_d \neq 0$, whereas a cumulant factor is *priced* only when it commands a nonzero premium in the cross section, an empirical question settled by the freely estimated Fama–MacBeth prices, not by the affine restriction. The restriction is refutable directly in levels: the freely estimated prices should collapse to the signed monomial $\eta_{p,q} = \prod_d b_d^{q_d}$, so the residual from fitting that monomial to the estimated prices is a direct collapse test.

The proposition also reframes the dimensionality question of the factor zoo. An affine D -factor SDF does *not* imply D priced factors: it implies $J(D, P)$ cumulant factors at every truncation order P , tied together by only D structural parameters. The “how many factors” question hence splits into two separate questions: how many factors span the SDF (D), and at what cumulant order the expansion can be truncated (P). A small D is compatible with a combinatorially rich set of cumulant factors, which is the structural sense in which the framework compresses the factor zoo.

The compression looks even sharper from the gross-return-space. The $J(D, P)$ count is the dimensionality in *cumulant* space; mapping back to gross returns multiplies it further. Because the gross-space structure is multiplicative, $\mathbb{E}[R_i] = R_f \prod_n \Lambda_n^{B_i^{(n)}}$, an SDF that is affine in D factors and truncated at order P generates in gross-space a web of interaction terms among the factors, with products of the cumulant factors rather than their sum. The apparent high dimensionality of the additive factor zoo is, on this reading, the gross-space shadow of a low-dimensional multiplicative object, itself governed by few structural parameters.

Table 2 fixes the terminology and separates the three counts that the “how many factors” question conflates: the span D , the theoretical cumulant-factor library $J(D, P)$, and the semi-aggregated subset actually priced in our reduced-form tests of Section 5.

Table 2. Risk factors, cumulant factors, and what is estimated

Object	Count at (D, P) = (6, 4)	Meaning
Risk factors (span)	$D = 6$	Observable variables f that span the SDF (Assumption 1); few in number.
Cumulant factors (library)	$J(D, P) = 116$	Joint-cumulant terms $Q_{i,p,q}$ of returns with the risk factors; the theoretical library generated at order P (Proposition 3), tied together by only D structural price coefficients.
Semi-aggregated factors (estimated)	$D + (P - 2) = 8$	Semi-aggregated terms actually priced in the Fama–MacBeth step: the D order-two directional terms plus one aggregate per higher order. A cumulant factor is <i>priced</i> when it earns a nonzero premium.

Cumulant-factor library, $\mathcal{B}(D, P)$, counts $J(D, P) = \sum_{h=1}^{P-1} (P-h) \binom{D+h-1}{h}$ terms:

	$P = 2$	$P = 3$	$P = 4$
$D = 1$	1	3	6
$D = 3$	3	12	31
$D = 6$	6	33	116

4.3. $\mathbb{K}$ -beta coordinates as estimable multivariate betas

A natural normalization of the latent-SDF exposures divides each co-cumulant by the SDF cumulant $\kappa_{a+b}(m)$, the within-order composite beta $B_i^{(n)}$ developed in Section 3.3. That normalization is the natural theoretical one, but it is not estimable: m is latent. With observable risk factors, there is an estimable normalization that generalizes the multivariate regression beta and keeps the structural prices interpretable.

Let $\Sigma_f \equiv \text{Var}(f)$ be the covariance matrix of the risk factors, assumed nonsingular, and $A \equiv \Sigma_f^{-1}$. The order-two case fixes the idea: a beta is a covariance divided by a variance, and with many factors, the multiple regression beta is $\Sigma_f^{-1} \text{Cov}(f, r_i)$; at order two, the coordinates defined here *are* the betas of standard empirical practice (in log-space). The general transform does the same thing at every order: apply Σ_f^{-1} to each factor argument of the joint cumulant. Formally, for each (p, h) group, the raw mixed cumulants over

ordered factor tuples form a symmetric order- h array $Z_i^{(p,h)}$ with entries

$$\left[Z_i^{(p,h)} \right]_{d_1 \dots d_h} = \kappa(\underbrace{r_i, \dots, r_i}_{p \text{ copies}}, f_{d_1}, \dots, f_{d_h}), \quad d_j \in \{1, \dots, D\},$$

and the $\mathbb{K}$ -beta transform applies A to every factor argument,

$$\widetilde{Z}_i^{(p,h)} = Z_i^{(p,h)} \times_1 A \times_2 \dots \times_h A$$

where $\times_k$ denotes the mode- k tensor product, multiplication of the k -th factor argument by A ; the formal tensor algebra is collected in Appendix A. The transformed array is again symmetric, so its entries depend on the index tuple only through the factor-copy pattern q ; write $\widetilde{Z}_{i,p,q}$ for that common value and $\widetilde{Q}_{i,p,q} \equiv c_{p,q} \widetilde{Z}_{i,p,q}$ for its eta-scaled version.

Proposition 4 ($\mathbb{K}$ -beta coordinates). *Let Σ_f be nonsingular and define $\theta \equiv \Sigma_f b$. Then:*

1. (Regression betas at order two.) *For $(p, h) = (1, 1)$, $\widetilde{Z}_i^{(1,1)} = \Sigma_f^{-1} \text{Cov}(f, r_i)$, the multiple regression beta vector of r_i on f .*
2. (Reparameterization.) *For each (p, h) , the map $Z_i^{(p,h)} \mapsto \widetilde{Z}_i^{(p,h)}$ is a linear bijection, with inverse given by applying Σ_f to every factor argument. $\mathbb{K}$ -beta coordinates and raw mixed cumulants carry the same pricing content.*
3. (Pricing in beta coordinates.) *Under Assumption 1,*

$$\mathbb{K}[r_i - r_f] = \sum_{p \geq 1, |q| \geq 1} \widetilde{\eta}_{p,q} \widetilde{Q}_{i,p,q}, \quad \widetilde{\eta}_{p,q} = \theta^q \equiv \prod_{d=1}^D \theta_d^{q_d}, \quad (19)$$

and $\theta = \text{Cov}(f, m)$: the structural price of every $\mathbb{K}$ -beta exposure is a monomial in the factor-level covariance prices $\theta_d = \text{Cov}(f_d, m)$.

Parts (1) and (2) say that the $\mathbb{K}$ -beta system is a reparameterization, not new information. It is the right reparameterization for estimation: the coordinates are regression-like objects, and the transform strips the scale and correlation structure of the risk factors out of the exposures, the exact multivariate generalization of dividing a covariance by a

variance. Part (3) says the structural prices in these coordinates are intrinsic economic objects. At order two, Equation (19) reads $\mathbb{K}[r_i - r_f]^{(2)} = \sum_d \theta_d \beta_{i,d}$, the log-space analogue of the textbook beta-pricing equation with covariance prices $\theta_d = \text{Cov}(f_d, m)$; at higher orders the same D numbers price everything, through products and powers θ^q . For $D = 1$, the coordinates reduce to $\frac{\kappa_{p,h}(r_i, f)}{\sigma^{2h}(f)}$ with prices $\text{Cov}(f, m)^h$. Every ingredient of Equation (19) except the prices is estimable from observables, where Σ_f and the joint cumulants $\kappa_{p,q}(r_i, f)$ require only return and factor data, so the $\mathbb{K}$ -beta coordinates are estimable multivariate betas, and the prices can be estimated freely by cross-sectional regression. That is precisely the monthly Fama–MacBeth design of Section 5.1.

Remark (Semi-aggregation is a one-price-per-order restriction). In the single-SDF theory of Sections 3.2 and 3.3, the within-order collapse to one composite beta and one price λ_n is *exact*, because every (a, b) term at total order n carries the same SDF cumulant $\kappa_n(m)$. With $D > 1$ risk factors, no exact collapse exists: by Proposition 4, terms at the same total order $n = p + |q|$ carry different structural prices θ^q whenever their factor-copy patterns differ. Summing the eta-scaled $\mathbb{K}$ -beta terms within total order and pricing each sum with a single coefficient (the empirical semi-aggregation of Section 5.1.3, which keeps the D order-two directional terms separate and aggregates orders three and above) is therefore a *restriction* on Equation (19), which trades a bias against the variance saved by estimating one price per order instead of one price per term.

With these results in hand, the reduced-form design of Section 5 is a sequence of theory-licensed steps. Proposition 2 licenses replacing the latent SDF with observable risk factors; Proposition 4 puts the observable exposures into estimable beta coordinates with interpretable structural prices; semi-aggregation imposes one price per order within the higher-order block; and the monthly Fama–MacBeth regressions estimate the prices freely, so the empirical tests assess the pricing content of the higher-order terms without imposing the structural restriction Equation (17), which remains available as an overidentifying diagnostic. Proofs of Propositions 2 to 4 are collected in Appendix A.

4.4. The overidentifying restrictions

The theory of this section culminates in a sharp, falsifiable prediction, which Section 5.3 takes to the data as the paper’s central empirical test. Under Assumption 1, the cross-sectional prices of the cumulant factors are not free parameters. By Proposition 4, the price of every $\mathbb{K}$ -beta exposure $\tilde{Q}_{i,p,q}$ is the signed monomial $\tilde{\eta}_{p,q} = \prod_d \theta_d^{q_d}$ in the D factor-level covariance prices $\theta_d = \text{Cov}(f_d, m)$. At order two, the model is just identified: D covariance prices for D betas, nothing to test. The overidentification begins at order three. At a given span D and truncation order P , the library contains $J(D, P)$ cumulant factors whose prices are pinned down by only D numbers, so the affine null imposes $J(D, P) - D$ overidentifying restrictions on the vector of freely estimated prices.

We test the restriction at the tractable span $D = 3$ and order $P = 4$: the $\mathbb{K}$ -Beta[3,4] model has $J(3, 4) = 31$ cumulant prices governed by three structural parameters, leaving *28 overidentifying restrictions*. The three-factor span is small enough that the 31 prices can be estimated freely, month by month, without the collinearity that makes the full 116-term six-factor system unstable, yet rich enough that the monomial structure has real content: the restriction forces the 27 higher-order prices to line up with products and powers of the three order-two prices, with no additional degrees of freedom. If the SDF is not affine in the chosen factors, so that higher-order risks carry premia untethered to the factor loadings, the freely estimated price vector will violate the product structure, and the restriction will be rejected. Section 5.3 implements this test for two three-factor spans, the Fama–French three factors, and (MKT, RMW, CMA), in price space and in return space, in sample and out of sample.

4.5. How many factors? Span versus order

A central question of the factor-zoo literature is dimensionality. [Cochrane \(2011\)](#) asks which characteristics provide independent information about average returns and how many factors are needed to summarize them; [Kozak et al. \(2020\)](#) find that no sparse

collection of named factors spans the SDF of a large anomaly panel, while a modest number of principal components does; [Bryzgalova et al. \(2023\)](#) average over the model space of linear factor combinations and conclude the SDF is dense in the candidate factors; [Kozak and Nagel \(2023\)](#) characterize when cross-sectional factors span the SDF at all, emphasizing that beta pricing with a given factor set is equivalent to the SDF being (conditionally) affine in those factors; and [Borri et al. \(2026\)](#) use a forward-selection Fama–MacBeth procedure to show that a few selected higher-order terms and interactions of common linear factors price the cross section and absorb much of the zoo. Common to all of these exercises is the object being counted: a factor is an additive term in a gross-space model, and dimensionality is the number of such terms with nonzero prices.

Proposition 3 says that, under exactly the affinity condition the spanning literature works with, this count conflates two distinct quantities. If the *log* pricing kernel (i.e. the SDF m) is affine in D risk factors, the objects carrying structural prices in the exact pricing identity are not D factors but the $J(D, P) = \sum_{h=1}^{P-1} (P-h) \binom{D+h-1}{h}$ cumulant factors of Equation (18) (116 at $D = 6, P = 4$), whose structural price coefficients are tied together by only D structural parameters through $\eta_{p,q} = \prod_d b_d^{q_d}$. A statistical census that counts priced regressors in an additive specification will therefore find “many factors” even when the SDF is genuinely low-dimensional, because every higher-order co-cumulant of returns with the D risk factors carries a nonzero structural price coefficient. The “how many factors” question splits into two separate questions with different answers: how many factors span the SDF (the span, D), and at what cumulant order the expansion can be truncated (the order, P).

This split offers a structural reading of the statistical findings. The principal components of a large anomaly panel are free to load on nonlinear transformations of a few underlying factors, so the [Kozak et al. \(2020\)](#) conclusion (no sparse factor model, but few PCs) is what a small- $D, P > 2$ economy looks like through a gross-return lens: sparsity fails in the space of named additive factors precisely because the cumulant-factor set is

the combinatorial expansion $J(D, P)$, while the PCs recover its low effective dimension. Likewise, the density of the SDF in [Bryzgalova et al. \(2023\)](#) is density in an additive candidate space; in cumulant space, the same SDF can be sparse in D and dense in the implied additive representation. And the affinity condition of [Kozak and Nagel \(2023\)](#) buys more than spanning at order two: in log space, it pins down the prices of the entire $J(D, P)$ expansion at once, which is what makes the question well-posed. The higher-order terms that [Borri et al. \(2026\)](#) select empirically are, in this reading, the gross-space shadow of the $J(D, P)$ cumulant factors: their forward-selection evidence that a few such terms absorb the zoo is the statistical counterpart to the structural compression derived here.

5. Empirics

Section [5.1](#) sets out the data and method; three questions then organize the evidence, one per subsection. First, do higher-order $\mathbb{K}$ -beta exposures improve out-of-sample pricing of the anomaly cross section relative to the raw factor benchmarks, and where in the cross-sectional distribution of pricing errors does any improvement live (Section [5.2](#))? Second, do the data reject the 28 overidentifying restrictions of Section [4.4](#), the structural prediction of the cumulant framework (Section [5.3](#))? Third, which ingredient of the model (the market factor, its higher cumulant orders, or the additional Fama–French factors) explains which anomaly (Section [5.4](#))? The answers, in brief: the pricing gains are real but concentrated in the *tails and significance* of the alpha distribution and in the Fama–MacBeth intercept, not in mean absolute alpha, where raw FF6 and the FF6-factor cumulant models are statistically indistinguishable; the overidentifying restrictions are not rejected in price space, and the restricted model matches or beats the unrestricted one in held-out fit; and the signal-level decomposition localizes the residual factor zoo to a handful of identifiable signals dominated by momentum-type persistence.

5.1. Data and method

We now take the cumulant decomposition to the data. Theorem 1 identifies the mixed co-cumulants of returns with the SDF as the carriers of priced risk and pins down their within-order weights in closed form. Our implementation is deliberately reduced-form. We replace the latent SDF with a small set of observable log risk factors, estimate the joint cumulants of each test asset with those factors, transform them into $\mathbb{K}$ -beta coordinates, compress them order by order in the spirit of the composite betas of Section 3.3, and price the resulting exposures in monthly Fama–MacBeth regressions in log space (Fama and MacBeth, 1973).

The bridge from theorem to implementation is the multilinearity of joint cumulants: if the SDF is conditionally affine in the chosen factors, every mixed co-cumulant $\kappa_{a,b}(r_i, m)$ in Equation (5) expands exactly into joint cumulants of the return with the risk factors, so the coordinates we price below are the estimable building blocks of the theorem’s expansion. The reduced form hence tests the pricing content of the higher-order terms without committing to a particular estimated SDF, and the pricing step is evaluated out of sample.

5.1.1. Test assets

Monthly portfolio returns are from the anomaly data of Chen and Zimmermann (2022) and span 1925-12 through 2024-12. The test asset universe, LS284, consists of the 284 long-short anomaly portfolios for which complete quintile information is available. For each signal, we form the collateralized long-short return: the simple return is the risk-free rate plus the spread between the economically long and economically short extreme quintiles, so the position is a fully collateralized overlay whose gross return is positive and whose log return is well defined. We treat these spreads as the primitive test assets: they are the objects the factor-zoo literature debates, and they put the burden of pricing on the long-minus-short dimension where anomaly alphas live. Throughout this section and the

next, $r_{i,t}^e \equiv r_{i,t} - r_{f,t}$ denotes the log excess return of the test asset i .

5.1.2. Risk factors

The risk factors are raw economic log factors, not standardized or whitened. The market factor is the log excess market return, $f_{M,t} = \log(R_{m,t}/R_{f,t})$. The long-short spread factor, $F_{d,t} \in \{\text{SMB}, \text{HML}, \text{RMW}, \text{CMA}, \text{UMD}\}$, is not itself a gross return, so we collateralize it the same way as the test assets: $f_{d,t} = \log((R_{f,t} + F_{d,t})/R_{f,t})$. We consider four factor sets: market only; the FF3 factors; the (MKT, RMW, CMA) set, hereafter MRC (Section 5.1.4); and the FF6 factors. Together with maximum cumulant orders, $P \in \{2, 3, 4\}$, this defines a ladder of twelve cumulant models, denoted $\mathbb{K}\text{-Beta}[D, P]$ for D risk factors and maximum order P , with the MRC variants labeled explicitly.

5.1.3. $\mathbb{K}$ -beta exposures

This subsection is the empirical projection of the theory of Section 4: it turns the $\mathbb{K}$ -beta coordinates of Proposition 4 into real-time estimates. For a model with D risk factors and maximum order P , the library contains every term indexed by $p \geq 1$ return copies and factor-copy pattern q with $1 \leq |q| \leq P - 1$ and $1 \leq p \leq P - |q|$. The raw mixed cumulant, $Z_{i,p,q,t} = \kappa_{p,q}(r_i^e, f)$, of asset i at month t is estimated from recursive exponentially weighted product moments using data through month t only, with cumulants recovered by the standard partition formula. The half-life is 180 months for total order two and 360 months for orders three and above; the minimum effective observation counts are 48, 120, and 180 for orders two, three, and four and above. (For the common-support comparison of Table 3, these minimum-history floors are harmonized across all sixteen models to 60, 90, and 120 months for orders two, three, and four and above, so that no model gains support from a looser floor.) Each term is scaled by its theory-implied coefficient $c_{p,q}$ of Proposition 2 and mapped into the $\mathbb{K}$ -beta coordinate $\tilde{Q}_{i,p,q,t}$ exactly as in Proposition 4, with the factor covariance matrix replaced by its the exponentially weighted moving

average (EWMA) estimate $\Sigma_{f,t}$.

Finally, we compress the transformed terms order by order, the empirical counterpart of the composite betas $B_i^{(n)}$. For multi-factor models, we preserve the D order-two directional terms separately and aggregate all order-three and order-four terms by total order. This imposes one price of risk per order within the higher-order block, a *restriction* on the unrestricted expansion, not an identity: for the richest model, the order-four six-factor $\mathbb{K}$ -beta model (labeled $\mathbb{K}$ -Beta[6, 4] in the tables and figures), it compresses the $J(6, 4) = 116$ cumulant factors of the theoretical library (Proposition 3) into eight semi-aggregated factors (six order-two directional terms plus one order-three and one order-four aggregate). The 116 is hence the theoretical count; the eight are what the cross-sectional regressions actually price.

5.1.4. Benchmark models

The raw CAPM, FF3, MRC, and FF6 benchmarks are estimated exactly as in standard practice: rolling 60-month time-series betas in simple excess-return space with at least 48 valid months, followed by monthly cross-sectional Fama–MacBeth regressions of simple excess returns on prior-month betas, intercept included.

MRC exists for a structural reason. The overidentification test of Section 5.3 needs a three-factor span, small enough that the 31 cumulant prices of the $\mathcal{B}(3, 4)$ library can be estimated freely month by month without the collinearity that makes the full 116-term six-factor system unstable, and it needs more than one such span, so that a non-rejection cannot be attributed to the idiosyncrasies of a single factor choice. FF3 is the canonical three-factor set; MRC replaces the size and value spreads with the profitability and investment spreads of Fama and French (2015), giving a second, economically distinct three-dimensional cut of the same factor family. Carrying the raw MRC benchmark and its cumulant ladder through the headline comparison also lets the reader see how this span prices the anomaly cross section before it is asked to carry the structural restriction.

Because the cumulant models forecast *log* excess returns, benchmark forecasts are mapped into log space with a direct *Jensen adjustment* estimated recursively from past data (namely, EWMA with a 180-month half-life and an effective minimum of roughly 68 months of history). The mapping is symmetric, and we verify in untabulated results that mirroring the empirical tests in simple-return rather than log-space (i.e., transforming the cumulants rather than the benchmarks via Jensen correction) leaves all rankings relevant to our claims intact and all numerical results nearly identical, so no model is advantaged by the choice of metric space.

5.1.5. Pricing and out-of-sample protocol

For each cumulant model and month $t + 1$, we run the cross-sectional regression

$$r_{i,t+1}^e = \lambda_{0,t+1} + X'_{i,t} \lambda_{t+1} + \epsilon_{i,t+1},$$

where $X_{i,t}$ stacks the model's $\mathbb{K}$ -beta exposures computed from data through month t . Exposures are never standardized, ranked, or whitened.

The out-of-sample period is 1978-07 through 2024-12. The forecast of month $t + 1$ applies the trailing 60-month mean of past monthly risk prices to the exposures dated t : $\widehat{\mu}_{i,t}^r = \bar{\lambda}_{0,t} + X'_{i,t} \bar{\lambda}_t$. Every ingredient of the forecast, including exposures, risk prices, and the Jensen adjustment for the raw benchmarks, uses only information available before the realized return being predicted.

5.1.6. Evaluation statistics

The primary pricing error is the out-of-sample log forecast error $\alpha_{i,t+1} = r_{i,t+1}^e - \widehat{\mu}_{i,t}^r$, and the per-asset alpha $\bar{\alpha}_i$ is its time average, reported in basis points per month. For each model, we report the cross-sectional mean absolute alpha (MAE), the distribution of per-asset alpha t -statistics computed using [Newey and West \(1987\)](#) standard errors with six lags, and the mean evaluated second-stage Fama–MacBeth intercept. The per-asset t -statistics are not a joint test on all assets together but rather describe the distribution of t per

asset; where formal inference on model comparisons is needed, we use a stationary block bootstrap over months (Politis and Romano, 1994; Politis, 2003) that resamples entire cross-sections to preserve contemporaneous error correlation in Section 5.2.1.

5.2. Headline results

The first empirical question is whether the higher-order $\mathbb{K}$ -beta exposures improve pricing out of sample: whether the cumulant ladders price the anomaly cross section better than the raw factor benchmarks, and, if so, where in the distribution of pricing errors the improvement lives.

5.2.1. Headline comparison

Table 3 reports the sixteen-model comparison on the common-support panel, our default baseline, in which every model is re-estimated with harmonized minimum-history requirements and evaluated on one identical 284-asset, 147,380 asset-month panel (1978-07 through 2024-12), so support-composition differences across models are removed by construction. Reading down the FF6-factor ladder (raw FF6, then $\mathbb{K}$ -Beta[6, 2] through $\mathbb{K}$ -Beta[6, 4]), three patterns stand out.

First, on the mean absolute alpha, the race is a statistical tie. Descriptively, the cumulant models edge raw FF6 (13.09–13.18 bps against 13.40 bps), but the block bootstrap of Section 5.1.6 shows the pairwise MAE differences between the FF6-factor cumulant models and raw FF6 are statistically indistinguishable from zero in either direction (for $\mathbb{K}$ -Beta[6, 4] versus raw FF6, -0.31 bps with a 95% confidence interval of $[-3.38, +2.23]$).

Second, alpha *significance* improves with cumulant order. Mean-squared alpha t -statistics fall monotonically from 1.47 under raw FF6 to 1.29, 1.27, and 1.26 as orders two through four are added; the share of assets with $|t_\alpha| > 1.96$ falls from 10.6% to 9.5%, and the share with $|t_\alpha| > 3$ from 2.1% to 1.4%.

Third, the mean Fama–MacBeth intercept (the average return the cross-sectional re-

Table 3. Out-of-sample log-alpha diagnostics on the common-support panel:
Raw benchmarks vs. $\mathbb{K}$ -Beta ladders (1978 – 2024)

Model	Exposures (#)	$ \bar{\alpha} $ MAE (log bps)	Mean t_{α}^2	$ t_{\alpha} > 1.96$ (% assets)	$ t_{\alpha} > 3$ (% assets)	FMB int. (bps)
<i>Panel A: Raw factor benchmarks</i>						
raw CAPM	1	15.36	1.92	13.0	4.9	21.35
raw FF3	3	14.26	1.74	12.3	3.9	20.78
raw FF6	6	13.40	1.47	10.6	2.1	16.89
raw (MKT,RMW,CMA)	3	13.29	1.53	10.6	3.5	19.36
<i>Panel B: $\mathbb{K}$-Beta cumulant ladders, grouped by factor span</i>						
<i>MKT only</i>						
$\mathbb{K}$ -Beta[1,2]	1	15.13	1.67	10.9	2.8	14.87
$\mathbb{K}$ -Beta[1,3]	2	14.93	1.64	11.6	2.5	14.38
$\mathbb{K}$ -Beta[1,4]	3	14.90	1.62	10.6	2.5	13.91
<i>FF3 span</i>						
$\mathbb{K}$ -Beta[3,2]	3	14.41	1.54	9.2	2.8	15.78
$\mathbb{K}$ -Beta[3,3]	4	14.38	1.49	9.9	1.8	15.06
$\mathbb{K}$ -Beta[3,4]	5	14.28	1.49	10.9	1.4	14.48
<i>MRC (MKT,RMW,CMA) span</i>						
$\mathbb{K}$ -Beta[3,2] (MRC)	3	13.72	1.41	7.0	2.5	13.90
$\mathbb{K}$ -Beta[3,3] (MRC)	4	13.63	1.39	7.0	1.8	13.35
$\mathbb{K}$ -Beta[3,4] (MRC)	5	13.45	1.32	7.0	1.8	12.77
<i>FF6 span</i>						
$\mathbb{K}$ -Beta[6,2]	6	13.18	1.29	8.8	1.8	12.14
$\mathbb{K}$ -Beta[6,3]	7	13.17	1.27	9.5	1.4	11.50
$\mathbb{K}$ -Beta[6,4]	8	13.09	1.26	9.5	1.4	11.05

Test assets are the LS284 collateralized long-short anomaly portfolios. Each model produces one-month-ahead out-of-sample forecasts of log excess returns using trailing 60-month mean Fama-MacBeth risk prices. $\mathbb{K}$ -Beta[D,P] denotes the $\mathbb{K}$ -beta cumulant model with risk-factor span D ($1 = \text{MKT}$; $3 = \text{FF3}$; $\text{MRC} = \text{MKT, RMW, CMA}$; $6 = \text{FF6}$) and maximum cumulant order P , with semi-aggregation by total order. All sixteen models are re-estimated with harmonized minimum-history requirements (60/90/120 months for orders 2/3/4+) and evaluated on one identical common-support panel of 284 assets and 147,380 asset-months (1978-07 through 2024-12). $|\bar{\alpha}|$ MAE is the cross-sectional mean absolute per-asset average forecast error in log bps per month. Mean t_{α}^2 and the $|t_{\alpha}|$ shares summarize the distribution of per-asset HAC alpha t -statistics (using the Newey-West standard error with six lags). FMB int. is 10^4 times the *signed* mean evaluated second-stage Fama-MacBeth intercept.

gressions cannot attribute to semi-aggregated factors) drops from 16.9 bps under raw FF6 to 11.5 bps under $\mathbb{K}$ -Beta[6, 3] and 11.1 bps under $\mathbb{K}$ -Beta[6, 4], a roughly one-third reduction in the level of unexplained average returns.

The market-only and intermediate ladders show the same qualitative pattern: significance improves with order even where MAE does not. Overall, the levels make clear that factor span matters; cumulant orders are a complement to, not a substitute for, the cross-sectional information in the six raw factors.

5.2.2. Tails of the alpha distribution

Table 4. Tail behavior of out-of-sample pricing errors across the LS284 cross section (common support, 1978 – 2024)

Model	$ \bar{\alpha}_i $ (bps/month)				$ t_{\alpha,i} $			Share (% assets)	
	p50	p90	p95	max	p50	p90	p95	$ t > 1.96$	$ t > 3$
raw CAPM	12.8	28.0	32.5	289	0.79	2.30	2.86	13.0	4.9
raw FF3	11.3	24.1	31.4	278	0.69	2.09	2.71	12.3	3.9
raw FF6	10.6	22.8	30.1	232	0.68	2.04	2.40	10.6	2.1
$\mathbb{K}$ -Beta[6,2]	9.6	25.6	33.5	192	0.67	1.83	2.23	8.8	1.8
$\mathbb{K}$ -Beta[6,3]	9.4	26.7	34.0	184	0.66	1.74	2.26	9.5	1.4
$\mathbb{K}$ -Beta[6,4]	9.2	26.0	35.0	186	0.65	1.82	2.23	9.5	1.4

Cross-sectional distribution of per-asset average out-of-sample log-alpha magnitudes $|\bar{\alpha}_i|$ (in bps per month) and of HAC alpha t -statistic magnitudes $|t_{\alpha,i}|$ across the LS284 long-short test assets, on each model's full evaluation panel. p50/p90/p95 denote percentiles; the last two columns report the share of assets whose alpha t -statistic exceeds 1.96 and 3 in absolute value. Per-asset t -statistics use the Newey-West standard errors with six lags and are descriptive distributional statistics, not a joint test. See the block-bootstrap inference summarized in Section 5.2.1.

Table 4 unpacks the distributional claim. In t -statistic space, the cumulant models improve on the benchmarks throughout the distribution: the median $|t_{\alpha,i}|$ falls from 0.68 (raw FF6) to 0.65 ($\mathbb{K}$ -Beta[6, 4]), the 90th percentile from 2.04 to 1.82, and the 95th percentile from 2.40 to 2.23. In raw alpha-magnitude space, the picture is mixed: the cumulant models price the typical asset better (median $|\bar{\alpha}_i|$ of 9.2–9.6 bps versus 10.6 bps for raw FF6), while their upper percentiles are somewhat fatter (95th percentile of 35.0 bps for $\mathbb{K}$ -Beta[6, 4] against 30.1 bps for raw FF6).

5.3. Testing the cumulant-price restrictions

We now test the paper’s central structural prediction, the 28 overidentifying restrictions of Section 4.4. The test operates on the $\mathbb{K}$ -Beta[3,4] model at two three-factor spans, the Fama–French three factors (FF3) and the MRC set of Section 5.1.4, and on a different cut of the same anomaly data from the long-short spreads of Section 5.1: the quintile-portfolio universe of $284 \text{ signals} \times 5 \text{ quintile bins}$, whose richer cross section (roughly 1,400 assets per month) is needed to estimate 31 free prices month by month. Exposures are the $\mathbb{K}$ -beta coordinates of Section 5.1.3, computed from the same recursive EWMA cumulant machinery.

The test compares two nested descriptions of the same cross section. The *unrestricted* model estimates 31 free prices of risk each month, one per cumulant factor; the *restricted* model generates the same 31 prices of risk from three loadings via the monomial map. The question is whether freeing the extra 28 parameters helps more than chance would predict, and we ask it in price space (Panel A of Table 5), in sample (Panel B), and out of sample (Panel C).

5.3.1. Test design

The building block is the unrestricted monthly price vector. For each month t , we run the cross-sectional regression of realized excess log returns on the lagged $\mathbb{K}$ -beta exposures, a standard Fama–MacBeth month with a wider regressor set,

$$r_{i,t}^e = a_t + z'_{i,t-1} \widehat{\eta}_t + e_{i,t}, \quad (20)$$

where $z_{i,t-1}$ stacks the 31 exposures of the $\mathcal{B}(3,4)$ library and $\widehat{\eta}_t \in \mathbb{R}^{31}$ is the freely estimated monthly price vector. The affine null says the same 31 prices are generated by three loadings, $\eta_{p,q}(b_t) = \prod_d b_{d,t}^{q_d}$. Each month, we fit this restricted model to the freely estimated

prices by weighted nonlinear least squares, a classical minimum-distance estimation,

$$\widehat{b}_t^W = \arg \min_{b \in \mathbb{R}^3} (\widehat{\eta}_t - \eta(b))' W_t (\widehat{\eta}_t - \eta(b)), \quad u_t^W \equiv \widehat{\eta}_t - \eta(\widehat{b}_t^W), \quad (21)$$

and collect the price residual u_t^W , the part of the monthly price vector that the restriction cannot reproduce.² The weight matrix W_t is diagonal, and we report two schemes: *pre-sample diagonal* weights, fixed from pre-1978 price dispersion, and *EWMA diagonal* weights, which track dispersion adaptively with a backward-looking 120-month half-life. The restriction is then evaluated on the time-averaged residual vector $\bar{u}^W = \frac{1}{T} \sum_t u_t^W$ through the standardized quadratic form

$$J = s(\bar{u}^W)' \widehat{\Omega}^{-1} s(\bar{u}^W), \quad (22)$$

where $s(\cdot)$ standardizes each coordinate by its bootstrap standard deviation and $\widehat{\Omega}$ is the (regularized) bootstrap correlation matrix of the standardized draws. This is the familiar GMM overidentification logic of Hansen (1982): fit the restricted model, form a quadratic form in what it cannot explain, and ask whether the residual exceeds the sampling noise, using a bootstrap reference distribution in place of the asymptotic χ^2 . No asymptotic distribution is reliable here, because the exposures are generated regressors estimated from overlapping EWMA windows, and the restricted-model fit is nonlinear. Both the standardization and the reference distribution of J therefore come from a circular block bootstrap over calendar months with a block length of 120, which preserves the slow-moving dependence induced by EWMA estimation.

Two reference distributions bracket the inference, and they answer different questions. The *empirical bootstrap* resamples the observed panel, reruns the entire pipeline (cumulant estimation, monthly regressions, restricted-model fits) on each pseudo-panel, and recenters the draws at the observed vector; it measures pure sampling dispersion, so its p -value asks whether $\bar{u}^W$ is far from zero relative to its own sampling noise. The *null-imposed*

² For readers who prefer the geometry: $\widehat{\eta}_t$ is a noisy point in 31 dimensions, the restricted model traces out a three-dimensional surface, and u_t^W is the gap to the nearest point on that surface.

bootstrap instead generates pseudo-panels in which the restriction holds by construction (returns are rebuilt from the restricted prices $\eta(\widehat{b}_t^W)$ plus block-resampled residual cross-sections) and reruns the full pipeline without recentering; its p -value is the test of the restrictions proper, asking whether the observed pipeline produces a larger aggregate price residual than a pipeline in which the restriction is true. Because the standardization in Equation (22) is reference-specific, the observed J differs across the two blocks by construction; J should be read as a standardized statistic, not an absolute quadratic form with χ_{28}^2 degrees-of-freedom intuition.

5.3.2. Results

Table 5 reports the evidence in three panels: price-space tests of the restrictions (Panel A), in-sample fit gains from relaxing them (Panel B), and held-out predictive comparisons (Panel C). All three panels share one inference convention: each observed statistic is compared with a bootstrap reference distribution, whose mean and standard deviation we report alongside it, and each p -value is the right-tail probability of a bootstrap draw at least as large as the observed statistic, computed from $B = 99$ circular block-bootstrap replications with a block length of 120 months.

For example, in the first row of Panel A, under the FF3 span with pre-sample weights, the observed price-residual statistic is $J = 33.17$; the empirical bootstrap regenerates the statistic with mean 40.40 and standard deviation 20.08; and the p -value of 0.63 records that 63% of the resampled draws exceed the observed value. A small p -value would signal a violation, an observed residual too large to attribute to estimation noise; a p -value near one says the observed residual sits inside, or below, what the reference machinery generates on its own.

Panel A shows that the restrictions are not rejected in price space. The null-imposed bootstrap is the test of the restrictions proper, and across both spans its p -values are 0.99–1.00: the observed aggregate price residual is smaller than what a pipeline in which the

Table 5. Testing the overidentifying restrictions of the $\mathbb{K}$ -Beta[3, 4] model

Panel A: Price-space overidentification tests

Factors	Weighting	Empirical bootstrap				Null-imposed bootstrap			
		J	mean	sd	p	J	mean	sd	p
FF3	Pre-sample diagonal	33.17	40.40	20.08	0.63	156.31	281.39	33.33	1.00
FF3	EWMA diagonal	35.68	39.09	19.89	0.53	140.35	251.69	30.91	1.00
MRC	Pre-sample diagonal	58.75	67.29	19.54	0.64	93.70	171.39	26.40	1.00
MRC	EWMA diagonal	60.72	62.44	19.50	0.48	120.09	167.02	24.99	0.99

Panel B: In-sample fit gains from relaxing the restriction

Factors	Restricted price path	Gain (%)	Empirical bootstrap			Return-space null bootstrap		
			mean	sd	p	mean	sd	p
FF3	Monthly restricted prices	36.97	36.99	3.22	0.40	37.02	2.18	0.45
FF3	Smoothed restricted prices	64.84	70.13	12.04	0.62	59.63	3.34	0.09
MRC	Monthly restricted prices	33.86	37.02	3.13	0.87	30.96	1.87	0.06
MRC	Smoothed restricted prices	59.13	69.29	13.67	0.70	71.21	16.33	0.67

Panel C: Predictive (held-out) fit comparison

Factors	Forecast rule	Gain (%)	Empirical bootstrap		
			mean	sd	p
FF3	EWMA (past-only) forecast	-0.68	-234.42	1147.25	0.17
FF3	Time-block crossfit	-0.73	-36.74	99.55	0.26
MRC	EWMA (past-only) forecast	-21.00	-4.93	9.95	0.94
MRC	Time-block crossfit	-5.95	-2.52	5.16	0.91

Tests of the 28 overidentifying restrictions of the $\mathbb{K}$ -Beta[3, 4] model (Section 5.3.1 describes the design) on the quintile-portfolio universe (284 signals $\times$ 5 bins). All reference distributions are circular block bootstraps (using a block length of 120 months and 99 replications; the MRC time-block crossfit row retains 95 valid replications), and all p -values are right-tail $(1 + \#\{\text{boot} \geq \text{obs}\})/(B + 1)$. Panel A: J is the bootstrap-standardized quadratic form in the mean price residual from the manifold projection; the null-imposed p is the test of the restrictions, and the two observed J entries differ because the standardization is reference-specific. Panel B: in-sample gain $100 \times (\text{SSE}_{\text{restr}} - \text{SSE}_{\text{unrestr}})/\text{SSE}_{\text{restr}}$ from freeing all 31 monthly prices. Panel C: the same gain on held-out data; negative values mean the restricted forecast fits better.

restriction holds by construction typically generates from estimation noise alone. The empirical-bootstrap p -values of 0.48–0.64 tell the same story from the other side: the mean price residual is well within its own sampling variation. Two caveats apply when reading Panel A. First, the two observed J entries in each row are not comparable in level (33.17 against 156.31 in the first row), because each is standardized by its own reference distribution; only the position of each J within its own block, summarized by the p -value, is informative. Second, the pre-sample and EWMA weighting rows are a robustness axis rather than separate tests, and they agree throughout. The non-rejection is unambiguous, but may also indicate that the estimation noise in 31 monthly prices is large relative to the deviations the test could detect.

Panel B shows that the in-sample cost of the restriction is what overfitting predicts. Freeing all 31 monthly prices rather than constraining them to the three-parameter structure improves in-sample return fit mechanically; the question is whether the observed gain exceeds what free parameters extract from noise, and the return-space null bootstrap, in which pseudo-panels are built so the restriction holds, makes that benchmark precise. Each span is evaluated against two restricted price paths: the *monthly* path re-estimates the three loadings freely each month, and the *smoothed* path adds a mild penalty on month-to-month movement in the loadings ($\rho \|b_t - b_{t-1}\|^2$ with $\rho = 1$), a stiffer benchmark that suppresses period-by-period repricing (the penalty affects estimation only; reported fit is the raw return residual in both cases). Against the monthly path, the unrestricted model fits 36.97% (FF3) and 33.86% (MRC) more of the return sum of squares, while the same relaxation harvests average gains of 37.0% and 31.0% from null panels with no restriction violation at all: free prices extract essentially the same fit gain from data in which the restriction is true. Against the smoothed path the gains are mechanically larger, 64.84% and 59.13%, because the stiffer benchmark fits worse to begin with, and the null benchmark rises nearly in step (59.6% and 71.2%). The right-tail p -values put the observed gains at or near the overfitting benchmark in all four cells: 0.45 and 0.67 in the two clear-cut cells,

and 0.06 (MRC, monthly) and 0.09 (FF3, smoothed) in the two marginal ones, which sit roughly one and a half null standard deviations above the overfitting mean and, at the resolution $B = 99$ replications affords, are suggestive rather than sharp. We conclude that in-sample fit gains of this magnitude do not provide evidence against the restriction.

Panel C asks the sharper, predictive question: if the 28 restrictions were binding constraints on true pricing, the restricted forecasts should lose on held-out data. They do not. Under both predictive designs, an EWMA rule using only past months and a five-fold time-block crossfit, the fit gain of the unrestricted model is negative: -0.68% and -0.73% for FF3, and -21.0% and -5.9% for MRC. The three-parameter restricted model forecasts held-out returns better than the 31-parameter unrestricted model, decisively so for the MRC span. One caveat applies to the bootstrap columns. The gain is a ratio of held-out sums of squares, and on resampled panels the 31-parameter forecast rule occasionally fails badly, so the reference distribution is heavy-tailed in gain units (mean -234.42 and standard deviation $1,147.25$ in the FF3 EWMA row); the mean and standard deviation columns describe that reference distribution, not alternative estimates of the observed gain, and the p -value, a rank within the distribution, is unaffected by its tails. The right-tail p -values of 0.17 – 0.94 confirm that no cell shows the unrestricted model outperforming beyond what resampling noise produces, and the instability of the 31-parameter rule on resampled panels is itself the variance cost the restriction removes. This pattern is consistent with the notion that the restriction acts as a shrinkage estimator whose target is dictated by the Euler equation rather than chosen statistically: if the theory is right, imposing the restriction is (approximately) a bias-free variance reduction, and the restricted model should win out of sample. The directional statement is consistent across all four cells: nothing is lost, and for MRC, much is gained by imposing the structure.

Overall, the three panels support the affine restriction in its economically relevant form. The freely estimated cumulant prices do not deviate detectably from the three-parameter restriction (Panel A); the in-sample fit gain the extra 28 parameters buy is indistinguishable

from overfitting (Panel B); and out of sample, the restricted model matches or beats the unrestricted one (Panel C).

5.4. Decomposition: Where the restriction bites

The tests so far treat the anomaly cross section as a whole. This subsection decomposes the same out-of-sample estimates family by family, asking two related questions. First, which Fama–French and higher-order factors characterize each family? Second, after those factors are priced, which families retain economically meaningful average forecast errors?

The mapping assigns the long-short signals to seventeen economically interpretable families. Per-signal assignment is reported in Internet Appendix [IA6](#). The first six rows are aligned with the FF6 factors: risk/volatility, size/age, value, profitability, investment, and momentum. The remaining rows collect earnings quality, other, distress, issuance, liquidity, intangibles, seasonality, accruals, options/volatility, reversal, and short interest. The plotted sample contains 283 signals for which the family-level exposure and pricing objects are jointly available. The row labels in Figures [1](#) and [2](#) report the number of retained signals in each family.

For a signal-level object $Y_{i,k}$, where k denotes a loading, higher-order exposure, mean return, or alpha, the family mean is

$$\bar{Y}_{g,k} = \frac{1}{N_g} \sum_{i:G(i)=g} \bar{Y}_{i,k}. \quad (23)$$

When $Y_{i,k,t}$ varies over time, $\bar{Y}_{i,k}$ is first formed as the signal-specific time-series mean over the dates on which that statistic is valid. We then average those signal-level means equally within the family. For the first six columns of Figure [1](#), $\bar{Y}_{i,k}$ is the mean of rolling FF6 loadings. The next two columns report the mean semi-aggregated order-three or order-four $\mathbb{K}$ -beta exposures. The pricing block contains three log-space objects. Let $r_{i,t+1}^e$ denote the realized log excess return, $\widehat{\mu}_{i,t}^{\text{FF6}}$ the one-month-ahead raw FF6 log forecast

after the direct Jensen adjustment, and $\widehat{\mu}_{i,t}^{K64}$ the $\mathbb{K}$ -Beta[6, 4] log forecast. The signal-level statistics are

$$\begin{aligned}\bar{r}_i &= 10^4 \times \frac{1}{T_i} \sum_{t \in \mathcal{T}_i} r_{i,t+1}^e, \\ \alpha_i^{\text{FF6}} &= 10^4 \times \frac{1}{T_i} \sum_{t \in \mathcal{T}_i} (r_{i,t+1}^e - \widehat{\mu}_{i,t}^{\text{FF6}}), \\ \alpha_i^{K64} &= 10^4 \times \frac{1}{T_i} \sum_{t \in \mathcal{T}_i} (r_{i,t+1}^e - \widehat{\mu}_{i,t}^{K64}).\end{aligned}$$

All three are in log basis points per month. The same signal-specific month set $\mathcal{T}_i$ is used for the three objects in the family display, so movement from the mean return to either alpha is not a support-composition effect.

Figure 1 standardizes the family means *within each column*. For $G = 17$ families,

$$Z_{g,k} = \frac{\bar{Y}_{g,k} - \mu_k}{s_k}, \quad \text{where } \mu_k = \frac{1}{G} \sum_{g=1}^G \bar{Y}_{g,k} \text{ and } s_k^2 = \frac{1}{G-1} \sum_{g=1}^G (\bar{Y}_{g,k} - \mu_k)^2. \quad (24)$$

Both the cell labels and the colors report $Z_{g,k}$. A red cell means that a family is high relative to other families in that column, and a blue cell means that it is low. The statistic is dimensionless. Neither a zero-pricing z-score nor a white-pricing cell means zero alpha; it means that the family lies near the cross-family average in that pricing column.

Figure 1 recovers the main economic structure expected from the family map. Five of the six canonical rows have a clear corresponding FF6 signature. Risk/volatility is high in the MKT column (1.79); value is the dominant HML family (2.90); profitability is the dominant RMW family (3.11); investment is high in CMA (1.76); and momentum is the dominant UMD family (2.80). These diagonal cells provide a direct check that the family labels summarize economically related signals rather than arbitrary collections of returns.

The size/age row is less clean. It contains only three signals, and its largest order-two signatures occur in RMW and CMA rather than exclusively in SMB. The row also has the most extreme order-four exposure signature (3.43). With $N = 3$, the family mean is sensitive to individual signal definitions, and the broad loading pattern should not be




Figure 1. Anomaly-family exposure heatmap.

Rows represent the seventeen signal families of 283 anomalies; the first six rows are the FF6-aligned families. The left block reports within-column z-scores of family mean FF6 factor loadings and aggregate Order-three and Order-four $\mathbb{K}$ -beta exposures (*loadings & exposures*). The right block reports within-column z-scores of out-of-sample pricing-error statistics (*pricing*), including the family mean log returns, raw FF6 alphas, and $\mathbb{K}$ -Beta[6, 4] alphas. Cell labels and colors are the same z-score so that within-column comparisons are visible across statistics on different scales. Colors are capped at $[-2, 2]$, while labels are uncapped.

treated as a population-level statement about all size-related anomalies.

The higher-order block reveals a structure that is not visible in the six ordinary FF loadings. Risk/volatility is high in Order 3 (1.42) and low in Order 4 (−1.02). Momentum has the strongest positive Order 3 exposure (2.18), while its Order 4 exposure is close to the cross-family mean. Liquidity and reversal have moderately positive Order 4 exposure, and options/volatility is low in Order 3 but positive in Order 4.

The pricing block summarizes the cross-family location of the three log-space outcomes. Short interest is the clear positive outlier in all three columns: its mean-return, FF6-alpha, and $\mathbb{K}$ -Beta-alpha z-scores are 2.61, 2.99, and 2.84. Options/volatility is the corresponding negative outlier, with z-scores of −2.12, −1.97, and −1.70. Momentum, seasonality, and reversal also lie above the cross-family average after both risk adjustments. Earnings quality, distress, and liquidity lie below it. These ranks are useful for locating the tails of the family distribution, but they are insufficient to evaluate whether a model prices a family. For that question, the relevant benchmark is zero alpha in basis points, which Figure 2 presents.

Figure 2 restores the baseline log-basis-point scale and makes zero the economic reference point. The distance from the hollow circle to an alpha marker is the average return absorbed by the corresponding model; the distance from the alpha marker to zero is the family-level residual. A negative alpha means that the model forecast exceeds the realized family mean on average, while a positive alpha means that the model underpredicts it. The figure is descriptive: a family mean alpha is not a test of joint significance, and small- N rows should be read as concentrated averages rather than broad regularities.

The canonical profitability and investment families are largely absorbed by the FF6 span. Profitability has a mean return of 30.3 log bps per month, a raw FF6 alpha of 2.6 bps, and a $\mathbb{K}$ -Beta[6,4] alpha of 3.1 bps. Investment moves from a mean return of 15.9 bps to an FF6 alpha of 0.2 bps and a $\mathbb{K}$ -Beta alpha of −1.4 bps. Issuance shows a similar pattern: its mean return is 17.7 bps, while its $\mathbb{K}$ -Beta alpha is only 0.1 bps. These rows are




Figure 2. Family-level mean returns and out-of-sample alphas in log space.

Hollow circles report family mean log excess returns, orange squares report raw FF6 alphas, and blue diamonds report $\mathbb{K}$ -Beta[6, 4] alphas. Each point is an equal-signal average of the corresponding signal-level statistic over the common support used for the family display. The dashed vertical line marks zero. Units are log basis points per month.

the cleanest examples of family premia that coincide with the economic directions the FF factors were designed to span.

Value is also substantially absorbed. Its mean return is 14.5 bps, compared with an FF6 alpha of 2.4 bps and a $\mathbb{K}$ -Beta alpha of 5.5 bps. Intangibles has a mean return of 9.4 bps; the $\mathbb{K}$ -Beta alpha is 3.5 bps, down from an FF6 alpha of 6.9 bps. Accruals improve more sharply, from an FF6 alpha of -6.6 bps to a $\mathbb{K}$ -Beta alpha of -3.0 bps. Earnings quality and the Other family also move closer to zero under $\mathbb{K}$ -Beta[6,4], from -8.2 to -6.9 bps and from -5.6 to -4.4 bps, respectively.

The remaining positive residuals are concentrated in persistence-related families. Momentum has a mean return of 32.5 bps and retains a $\mathbb{K}$ -Beta alpha of 9.9 bps. Seasonality retains 13.0 bps, and reversal retains 13.8 bps. Short interest is the dominant residual family: its mean return is 62.0 bps, its raw FF6 alpha is 49.0 bps, and its $\mathbb{K}$ -Beta alpha remains 43.8 bps. The $\mathbb{K}$ -Beta model reduces the short-interest residual by about 5.2 bps relative to FF6, but most of the family premium remains.

Several families retain negative average pricing errors. Risk/volatility has a $\mathbb{K}$ -Beta alpha of -9.1 bps, earnings quality -6.9 bps, distress -9.7 bps, liquidity -6.1 bps, and options/volatility -25.6 bps. The options/volatility row is unusual because its mean return is already negative (-24.3 bps); both models leave a negative alpha of similar magnitude. Size/age also moves materially away from zero under $\mathbb{K}$ -Beta[6,4], from an FF6 alpha of -5.4 bps to -19.2 bps. The small number of signals and the row's unusual exposures in Figure 1 make this an example where the family average should be treated as a diagnostic.

The two figures are complementary. Several rows illustrate why both are needed. Profitability and investment have the expected RMW and CMA signatures and are nearly priced in Figure 2. Momentum has the expected UMD signature and a large Order 3 exposure, yet retains a positive alpha near 10 bps; the family is exposed to included directions, but those directions do not absorb the full premium. Short interest is extreme in the pricing block but does not have a comparably simple FF6 loading signature, con-

sistent with a residual that is not explained by a single included direction. Conversely, size/age has unusually large exposure across several columns but a negative $\mathbb{K}$ -Beta alpha, demonstrating that unusual exposures do not, by themselves, imply a positive risk premium.

Table 6. The 18 survivors: Signals the $\mathbb{K}$ -Beta[6, 4] model does not price

Signal	Family	Mean ret		raw FF3		raw FF6		$\mathbb{K}$ -Beta[6, 4]	
		$\bar{\alpha}$	t	$\bar{\alpha}$	t	$\bar{\alpha}$	t	$\bar{\alpha}$	t
IndRetBig	momentum	137	7.62	120	6.84	122	6.96	107	5.03
MomSeason06YrPlus	seasonality	49	5.42	33	3.68	38	4.16	41	3.79
ReturnSkew3F	risk vol	-10	-1.72	-29	-4.88	-24	-3.88	-24	-3.39
TrendFactor	momentum	87	5.94	70	4.97	63	4.38	52	3.06
pchquick	distress	2	0.46	-19	-3.54	-17	-2.94	-18	-2.95
Tax_q	earnings quality	4	0.53	-16	-2.04	-19	-2.36	-24	-2.83
IO_ShortInterest	short interest	250	4.23	285	4.80	234	3.94	192	2.82
NOA	investment	37	4.02	16	1.73	19	1.99	29	2.71
FirmAgeMom	momentum	93	3.82	71	2.87	55	2.19	75	2.59
MeanRankRevGrowth	investment	1	0.07	-15	-1.23	-15	-1.21	-31	-2.30
EarningsPredictability	earnings quality	44	3.20	23	1.74	28	2.00	35	2.25
MomSeason16YrPlus	seasonality	27	3.12	10	1.12	16	1.77	23	2.24
GrSaleToGrOverhead	other	-5	-0.52	-25	-2.79	-20	-2.18	-22	-2.18
DelLTI	issuance	7	1.41	-16	-3.00	-11	-2.11	-12	-2.09
HerfBE	intangibles	-3	-0.36	-23	-2.98	-19	-2.41	-18	-2.04
OperProfRDLAGAT_q	profitability	65	4.33	49	3.25	43	2.86	36	2.03
MomSeason	seasonality	34	2.77	14	1.16	17	1.39	28	2.00
SmileSlope	options vol	54	3.69	37	2.55	41	2.77	36	1.99

The 18 long-short signals whose $\mathbb{K}$ -Beta[6, 4] out-of-sample alpha is significant ($|t_\alpha| > 1.96$), ordered by $|t_\alpha|$ under $\mathbb{K}$ -Beta[6, 4]. Each pair of columns reports the per-signal average out-of-sample alpha $\bar{\alpha}$ (bps/month, signed) and its HAC t -statistic (using the Newey-West standard errors with six lags) under the corresponding model. “Mean ret” is the no-model baseline: the forecast is zero, so the alpha is the signal’s mean out-of-sample log excess return. Families follow the keyword classifier of Figure 1. A signal can be significant under $\mathbb{K}$ -Beta[6, 4] without a significant mean return and vice versa: pricing exposures change both the point estimate and its precision.

Table 6 then lists the survivors: the eighteen signals whose $\mathbb{K}$ -Beta[6, 4] out-of-sample alpha remains significant ($|t_\alpha| > 1.96$), with the mean-return baseline and the raw FF3 and FF6 alphas alongside for comparison. The list has three notable features. First, the survivors are concentrated in momentum- and trend-related signals: six of the eighteen (the industry return of big firms, the trend factor, firm-age momentum, and three seasonal-momentum signals) are momentum- or seasonality-type, the largest single block, consis-

tent with the introduction’s characterization; the rest scatter thinly across the remaining families. Second, the survivors are not an artifact of the cumulant machinery: raw FF6 fails on fifteen of the same eighteen signals and raw FF3 on thirteen, and the largest premia (short interest in particular, cf. Figure 2) are large under every model. Third, only eleven of the eighteen have a significant raw mean return to begin with; for the other seven, the premium is insignificant gross of any adjustment, and it is the risk adjustment itself that reveals a significant (mostly negative) residual, a reminder that pricing exposures moves both the point estimate and its precision. Overall, the decomposition localizes the residual factor zoo: after pricing six factors and their co-cumulants through order four, the unexplained cross section is not a diffuse zoo but a short, named list concentrated in momentum-type persistence (risk factors outside the FF6 span) rather than in missing higher orders of the included factors.

6. Conclusion

We write the Euler equation in cumulant space and obtain three structural results. Within each total order n , all mixed co-cumulants of asset returns and the SDF collapse into a single composite beta $B_i^{(n)}$ and a single price of risk λ_n , delivering a one-beta-per-order representation. The implied factor structure for gross expected returns is multiplicative, $\mathbb{E}[R_i] = R_f \prod_n \Lambda_n^{B_i^{(n)}}$, rather than additive. And the underlying log-space identity is an exact Euler-equation equivalence, a formalization and cross-sectional extension of the coentropy concept of Backus et al. (2018). When the SDF is affine in observable risk factors, the identity decomposes exactly onto those factors, the prices of the full $J(D, P)$ -term cumulant expansion collapse to monomials in the D SDF loadings, and the $\mathbb{K}$ -beta transform renders the exposures estimable.

The affine case does more than license estimation: it overidentifies the model. At a three-factor span and fourth cumulant order, the affine null ties 31 cumulant prices to three structural parameters, and the resulting 28 overidentifying restrictions are not rejected:

the freely estimated price vector deviates less from the three-parameter restriction than a pipeline in which the restriction holds by construction typically does, the in-sample fit gain from freeing all 31 prices is what overfitting alone produces, and on held-out data the restricted model forecasts as well as or better than the unrestricted one. Alongside, pricing the order-level $\mathbb{K}$ -beta exposures of 284 long-short anomaly portfolios out of sample shrinks the mean-squared alpha t -statistic from 1.47 under the Fama–French six-factor benchmark to 1.26 at the fourth order and reduces the mean Fama–MacBeth intercept by roughly one third, while mean absolute alphas remain statistically indistinguishable from the benchmark’s. These results show that a small number of structural parameters organizes the priced risk in a large anomaly cross section, and that the data do not demand more.

More generally, our results change how the factor zoo should be read. The gross-return factor structure implied by the law of one price is multiplicative, and the familiar additive factor models of gross-space are best understood as benchmark approximations of a smaller set of underlying cumulant factors. The dimensionality of priced risk splits into the span of the SDF and the order of the cumulant expansion: a small number of risk factors implies a combinatorially rich set of cumulant factors, which is why additive gross-return censuses of the zoo find many factors. In the data, after six factors and their co-cumulants through order four are priced, the unexplained residual consists of eighteen specific signals, concentrated in momentum-type persistence. The immediate next empirical question is therefore not whether ever-higher orders can be written down, but whether the second-, third-, and fourth-order terms of a few factors already provide an economically sufficient description of priced risk in the anomaly cross section, a question the non-rejection of the overidentifying restrictions answers in the affirmative at the tested three-factor spans. Richer spans and different risk factors can now be tested with the same machinery.

AI Disclosure Statement

During the preparation of this work, the authors used agentic coding assistants (Codex, Claude Code) for coding and editing. The author(s) reviewed and edited the output as needed and take full responsibility for the content of the published article.

References

- Tobias Adrian, Erkkko Etula, and Tyler Muir. Financial intermediaries and the cross-section of asset returns. *Journal of Finance*, 69(6):2557–2596, 2014.
- Fernando Alvarez and Urban J. Jermann. Using asset prices to measure the persistence of the marginal utility of wealth. *Econometrica*, 73(6):1977–2016, 2005.
- David Backus, Mikhail Chernov, and Ian Martin. Disasters implied by equity index options. *Journal of Finance*, 66(6):1969–2012, 2011.
- David Backus, Mikhail Chernov, and Stanley Zin. Sources of entropy in representative agent models. *Journal of Finance*, 69(1):51–99, 2014.
- David Backus, Nina Boyarchenko, and Mikhail Chernov. Term structures of asset prices and returns. *Journal of Financial Economics*, 129(1):1–23, 2018.
- Gurdip Bakshi, Nikunj Kapadia, and Dilip Madan. Stock return characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies*, 16(1):101–143, 2003.
- Ravi Bansal and Amir Yaron. Risks for the long run: A potential resolution of asset pricing puzzles. *Journal of Finance*, 59(4):1481–1509, 2004.
- Geert Bekaert and Eric Engstrom. Asset return dynamics under habits and bad environment-good environment fundamentals. *Journal of Political Economy*, 125(3):713–760, 2017.
- Fischer Black. Capital market equilibrium with restricted borrowing. *Journal of Business*, 45(3):444–455, 1972.
- Nicola Borri, Denis Chetverikov, Yukun Liu, and Aleh Tsyvinski. Higher-order asset

- pricing factors via forward selection Fama-MacBeth regression. SSRN Working Paper No. 5181447, 2026.
- Douglas T. Breeden. An intertemporal asset pricing model with stochastic consumption and investment opportunities. *Journal of Financial Economics*, 7(3):265–296, 1979.
- Svetlana Bryzgalova, Jiantao Huang, and Christian Julliard. Bayesian solutions for the factor zoo: We just ran two quadrillion models. *Journal of Finance*, 78(1):487–557, 2023.
- Mark M. Carhart. On persistence in mutual fund performance. *Journal of Finance*, 52(1):57–82, 1997.
- Fousseni Chabi-Yo. Pricing kernels with stochastic skewness and volatility risk. *Management Science*, 58(3):624–640, 2012.
- Fousseni Chabi-Yo and Riccardo Colacito. The term structures of coentropy in international financial markets. *Management Science*, 65(8):3541–3558, 2019.
- David A. Chapman. Approximating the asset pricing kernel. *Journal of Finance*, 52(4):1383–1410, 1997.
- Andrew Y. Chen and Tom Zimmermann. Open source cross-sectional asset pricing. *Critical Finance Review*, 11(2):207–264, 2022.
- John H. Cochrane. *Asset Pricing*. Princeton University Press, 2005.
- John H. Cochrane. Presidential address: Discount rates. *Journal of Finance*, 66(4):1047–1108, 2011.
- Robert F. Dittmar. Nonlinear pricing kernels, kurtosis preference, and evidence from the cross section of equity returns. *Journal of Finance*, 57(1):369–403, 2002.
- Eugene F. Fama and Kenneth R. French. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56, 1993.

- Eugene F. Fama and Kenneth R. French. A five-factor asset pricing model. *Journal of Financial Economics*, 116(1):1–22, 2015.
- Eugene F. Fama and James D. MacBeth. Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, 81(3):607–636, 1973.
- Guanhao Feng, Stefano Giglio, and Dacheng Xiu. Taming the factor zoo: A test of new factors. *Journal of Finance*, 75(3):1327–1370, 2020.
- Anisha Ghosh, Christian Julliard, and Alex P. Taylor. What is the consumption-CAPM missing? An information-theoretic framework for the analysis of asset pricing models. *Review of Financial Studies*, 30(2):442–504, 2017.
- Niels Joachim Gormsen and Christian Skov Jensen. Higher-moment risk. *Journal of Finance*, forthcoming, 2025.
- Shihao Gu, Bryan T. Kelly, and Dacheng Xiu. Empirical asset pricing via machine learning. *Review of Financial Studies*, 33(5):2223–2273, 2020.
- Lars Peter Hansen. Large sample properties of generalized method of moments estimators. *Econometrica*, 50(4):1029–1054, 1982.
- Lars Peter Hansen. Dynamic valuation decomposition within stochastic economies. *Econometrica*, 80(3):911–967, 2012.
- Lars Peter Hansen and Ravi Jagannathan. Implications of security market data for models of dynamic economies. *Journal of Political Economy*, 99(2):225–262, 1991.
- Lars Peter Hansen and Scott F. Richard. The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models. *Econometrica*, 55(3):587–613, 1987.
- Lars Peter Hansen and José A. Scheinkman. Long-term risk: An operator approach. *Econometrica*, 77(1):177–234, 2009.

- Campbell R. Harvey and Akhtar Siddique. Conditional skewness in asset pricing tests. *Journal of Finance*, 55(3):1263–1295, 2000.
- Campbell R. Harvey, Yan Liu, and Heqing Zhu. ...and the cross-section of expected returns. *Review of Financial Studies*, 29(1):5–68, 2016.
- Zhiguo He and Arvind Krishnamurthy. Intermediary asset pricing. *American Economic Review*, 103(2):732–770, 2013.
- Kewei Hou, Chen Xue, and Lu Zhang. Digesting anomalies: An investment approach. *Review of Financial Studies*, 28(3):650–705, 2015.
- Bryan T. Kelly, Seth Pruitt, and Yinan Su. Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics*, 134(3):501–524, 2019.
- Serhiy Kozak and Stefan Nagel. When do cross-sectional asset pricing factors span the stochastic discount factor? NBER Working Paper 31275, 2023.
- Serhiy Kozak, Stefan Nagel, and Sridhar Santosh. Shrinking the cross section. *Journal of Financial Economics*, 135(2):271–292, 2020.
- Alan Kraus and Robert H. Litzenberger. Skewness preference and the valuation of risk assets. *Journal of Finance*, 31(4):1085–1100, 1976.
- Albert S. Kyle and Karamfil Todorov. The cumulant risk premium. BIS Working Paper No. 1128, 2023.
- John Lintner. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics*, 47(1):13–37, 1965.
- Robert E. Lucas, Jr. Asset prices in an exchange economy. *Econometrica*, 46(6):1429–1445, 1978.

- Ian Martin. Consumption-based asset pricing with higher cumulants. *Review of Economic Studies*, 80(2):745–773, 2013.
- Ian Martin. What is the expected return on the market? *Quarterly Journal of Economics*, 132(1):367–433, 2017.
- Ian W. R. Martin and Ran Shi. On the moments of the stochastic discount factor. CEPR Discussion Paper No. 21235, 2026.
- Whitney K. Newey and Kenneth D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987.
- Piotr Orłowski, Andras Sali, and Fabio Trojani. Arbitrage free dispersion. Swiss Finance Institute Research Paper No. 19-20, 2019.
- Dimitris N. Politis. The impact of bootstrap methods on time series analysis. *Statistical Science*, 18(2):219–230, 2003.
- Dimitris N. Politis and Joseph P. Romano. The stationary bootstrap. *Journal of the American Statistical Association*, 89(428):1303–1313, 1994.
- Stephen A. Ross. The arbitrage theory of capital asset pricing. *Journal of Economic Theory*, 13(3):341–360, 1976.
- Paul Schneider and Fabio Trojani. (Almost) model-free recovery. *Journal of Finance*, 74(1):323–370, 2019.
- William F. Sharpe. Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, 19(3):425–442, 1964.
- Karl N. Snow. Diagnosing asset pricing models using the distribution of asset returns. *Journal of Finance*, 46(3):955–983, 1991.

Appendix

A. Proofs

Proof of Theorem 1

Statements (1) through (3) are equivalent by substitution: $M = e^{-m}$ and $R_i = e^{r_i}$ imply

$$\mathbb{E}[MR_i] = 1 \iff \mathbb{E}[e^{r_i - m}] = 1 \iff \mathbb{K}[r_i - m] = 0.$$

For statement (4), we expand the CGF at argument one:

$$\mathbb{K}[r_i - m] = \sum_{n=1}^{\infty} \frac{1}{n!} \kappa_n(r_i - m).$$

Hence, statement (3) is equivalent to statement (4).

For statement (5), we use the joint CGF of (r_i, m) :

$$K_{r_i, m}(s_1, s_2) \equiv \log \mathbb{E}[e^{s_1 r_i + s_2 m}] = \sum_{a, b \geq 0} \frac{s_1^a s_2^b}{a! b!} \kappa_{a, b}(r_i, m).$$

Evaluating at $(s_1, s_2) = (1, -1)$ yields

$$\mathbb{K}[r_i - m] = \sum_{a, b \geq 0} \frac{(-1)^b}{a! b!} \kappa_{a, b}(r_i, m),$$

which proves the equivalence of statements (3), (4), and (5).

To obtain statement (6), note that r_f is (conditionally) deterministic, so

$$\mathbb{K}[r_i - r_f] = \mathbb{K}[r_i] - r_f.$$

The risk-free Euler equation implies

$$0 = \mathbb{K}[r_f - m] = r_f + \mathbb{K}[-m], \quad r_f = -\mathbb{K}[-m].$$

Therefore,

$$\mathbb{K}[r_i - r_f] = \mathbb{K}[r_i] + \mathbb{K}[-m].$$

Now, let's separate the expansion of $\mathbb{K}[r_i - m] = 0$ in statement (5) into its pure-return block ($b = 0, a \geq 1$), its pure-SDF block ($a = 0, b \geq 1$), and its mixed block ($a, b \geq 1$). Because $\kappa_{a,0}(r_i, m) = \kappa_a(r_i)$ and $\kappa_{0,b}(r_i, m) = \kappa_b(m)$, the pure blocks sum to

$$\sum_{a \geq 1} \frac{\kappa_a(r_i)}{a!} + \sum_{b \geq 1} \frac{(-1)^b \kappa_b(m)}{b!} = \mathbb{K}[r_i] + \mathbb{K}[-m] = \mathbb{K}[r_i - r_f].$$

Moving the mixed block to the other side of $\mathbb{K}[r_i - m] = 0$ gives

$$\mathbb{K}[r_i - r_f] = \sum_{a,b \geq 1} \frac{(-1)^{b+1}}{a! b!} \kappa_{a,b}(r_i, m),$$

which is statement (6).

Statement (7) is a notational relabeling of statement (6) using $\beta_i^{(a,b)}$ and λ_n .

Finally, statement (8) follows from exponentiating the definition of $\mathbb{K}[r_i - r_f]$:

$$e^{\mathbb{K}[r_i - r_f]} = \mathbb{E}\left[\frac{R_i}{R_f}\right], \quad \mathbb{E}[R_i] = R_f e^{\mathbb{K}[r_i - r_f]}.$$

Substituting the order-level representation for $\mathbb{K}[r_i - r_f]$ gives

$$\mathbb{E}[R_i] = R_f \exp\left(\sum_{n \geq 2} B_i^{(n)} \lambda_n\right) = R_f \prod_{n \geq 2} \Lambda_n^{B_i^{(n)}},$$

which is statement (8). Each statement hence implies the others. □

Proof of Proposition 2

Joint cumulants of total order at least two are multilinear in each argument and invariant to adding a constant to any argument. Write $\kappa_{p,h}(r_i, m) = \kappa(r_i, \dots, r_i, m, \dots, m)$ with p return slots and h SDF slots, and substitute $m = a + \sum_d b_d f_d$ into each SDF slot. Expanding slot by slot, the constant a contributes zero in every slot because the total order $p + h$ is at least two, leaving the ordered-tuple expansion

$$\kappa_{p,h}(r_i, m) = \sum_{d_1=1}^D \cdots \sum_{d_h=1}^D \left(\prod_{j=1}^h b_{d_j} \right) \underbrace{\kappa(r_i, \dots, r_i, f_{d_1}, \dots, f_{d_h})}_{p \text{ copies}}. \quad (\text{A1})$$

Joint cumulants are symmetric in their arguments, so each summand in Equation (A1)

depends on the tuple $(d_1, \dots, d_h)$ only through its multiset, that is, through the multi-index q with $q_d = \#\{j : d_j = d\}$. The number of ordered tuples realizing a given q with $|q| = h$ is the multinomial coefficient $h! / \prod_d q_d!$, which yields Equation (15).

Substituting Equation (15) into statement (6) of Theorem 1, the term of return order p and SDF order h contributes

$$\frac{(-1)^{h+1}}{p! h!} \kappa_{p,h}(r_i, m) = \sum_{|q|=h} \frac{(-1)^{|q|+1}}{p! \prod_d q_d!} b^q \kappa_{p,q}(r_i, f) = \sum_{|q|=h} \eta_{p,q} Q_{i,p,q},$$

and summing over $p \geq 1$ and $h \geq 1$ gives Equation (16). $\square$

Proof of Proposition 3

That Equation (16) holds with $\eta_{p,q} = b^q$ is the second display of Proposition 2. Part (1) is immediate because b^q does not involve p . For part (2), fix the number of factor copies $h = |q| \in \{1, \dots, P-1\}$. The number of multi-indices $q \geq 0$ with $|q| = h$ is the number of weak compositions of h into D parts, $\binom{D+h-1}{h}$, and for each such q the return order p ranges over $\{1, \dots, P-h\}$, contributing the factor $(P-h)$. Summing over h gives Equation (18). $\square$

Remark (Identification). If the exposure vectors $(Q_{i,p,q})_{(p,q) \in \mathcal{B}(D,P)}$ span $\mathbb{R}^{J(D,P)}$ across assets, then any price vector satisfying the truncated pricing equation (with the residual $\zeta_i^{(P)}$ assigned to orders above P) is unique and therefore equals (b^q) ; without cross-sectional spanning the restriction still holds but is not separately identified.

Proof of Proposition 4

For tensors $U, V \in (\mathbb{R}^D)^{\otimes h}$, write $\langle U, V \rangle$ for the entrywise (Frobenius) inner product, and for $u \in \mathbb{R}^D$ write $u^{\otimes h}$ for the rank-one tensor with entries $\prod_{j=1}^h u_{d_j}$.

(i) For $(p, h) = (1, 1)$, the tensor $Z_i^{(1,1)}$ is the D -vector with entries $\kappa(r_i, f_d) = \text{Cov}(f_d, r_i)$, so $\widetilde{Z}_i^{(1,1)} = \Sigma_f^{-1} \text{Cov}(f, r_i)$, the slope vector of the L^2 projection of r_i on $(1, f)'$, the multivariate regression beta.

(ii) In flattened form, the mode-wise transform is $\text{vec}(\widetilde{Z}_i^{(p,h)}) = A^{\otimes h} \text{vec}(Z_i^{(p,h)})$. Since

$A = \Sigma_f^{-1}$ is invertible, $A^{\otimes h}$ is invertible with inverse $\Sigma_f^{\otimes h}$; that is, the inverse map applies Σ_f to every direction slot.

(iii) By the ordered-tuple expansion Equation (A1), the contribution of the (p, h) group to statement (6) of Theorem 1 is

$$\frac{(-1)^{h+1}}{p! h!} \langle b^{\otimes h}, Z_i^{(p, h)} \rangle.$$

Because A is symmetric and $\theta = \Sigma_f b$ satisfies $A\theta = b$,

$$\langle \theta^{\otimes h}, \widetilde{Z}_i^{(p, h)} \rangle = \langle \theta^{\otimes h}, A^{\otimes h} Z_i^{(p, h)} \rangle = \langle (A\theta)^{\otimes h}, Z_i^{(p, h)} \rangle = \langle b^{\otimes h}, Z_i^{(p, h)} \rangle.$$

The transformed tensor $\widetilde{Z}_i^{(p, h)}$ is symmetric, because $Z_i^{(p, h)}$ is symmetric and the same operator is applied to every slot, so collecting the ordered tuples in $\langle \theta^{\otimes h}, \widetilde{Z}_i^{(p, h)} \rangle$ into multisets gives

$$\frac{(-1)^{h+1}}{p! h!} \langle \theta^{\otimes h}, \widetilde{Z}_i^{(p, h)} \rangle = \sum_{|q|=h} c_{p, q} \theta^q \widetilde{Z}_{i, p, q} = \sum_{|q|=h} \widetilde{\eta}_{p, q} \widetilde{Q}_{i, p, q}.$$

Summing over $p \geq 1$ and $h \geq 1$ yields Equation (19). Finally, under Assumption 1, $\text{Cov}(f_d, m) = \text{Cov}(f_d, a + b' f) = (\Sigma_f b)_d = \theta_d$, since covariances are invariant to the additive constant a . □

Internet Appendix

IA1. Supporting Theory

IA1.1. From $\mathbb{K}$ -expected returns to expected returns

The pricing identity in Theorem 1 links $\mathbb{K}[r_i - r_f]$ to co-cumulants of returns and the SDF. Two exact algebraic routes convert that identity into estimating equations whose left-hand sides are conventional expected returns.

Route 1: expected excess log returns via the CGF. For any random variable X ,

$$\mathbb{K}[X] = \sum_{n=1}^{\infty} \frac{\kappa_n(X)}{n!}.$$

Setting $X = r_i - r_f$ and separating the first cumulant gives

$$\mathbb{K}[r_i - r_f] = \mu_i + \Omega_i, \quad \mu_i \equiv \mathbb{E}[r_i - r_f], \quad \Omega_i \equiv \sum_{n=2}^{\infty} \frac{\kappa_n(r_i)}{n!}. \quad (\text{IA1})$$

The object Ω_i is the own-risk correction, which converts $\mathbb{K}$ -expected returns into expected log returns. Since r_f is $\mathcal{F}_{t-1}$ -measurable, cumulants of order two and above are unaffected by subtracting it.

Substituting the order-level pricing rule into Equation (IA1) yields the observation equation

$$r_{i,t} - r_{f,t} = \sum_{n=2}^N \lambda_{n,t} B_{i,t}^{(n)} - \Omega_{i,t}^{(N)} + \epsilon_{i,t}, \quad \mathbb{E}_{t-1}[\epsilon_{i,t}] = 0. \quad (\text{IA2})$$

Route 2: expected excess simple returns via exact exponentiation. Exponentiating the definition of $\mathbb{K}[r_i - r_f]$ gives

$$e^{\mathbb{K}[r_i - r_f]} = \mathbb{E}\left[\frac{R_i}{R_f}\right].$$

Therefore

$$\mathbb{E}[R_i - R_f] = R_f (e^{\mathbb{K}[r_i - r_f]} - 1), \quad (\text{IA3})$$

which is exact. Substituting the order-level pricing rule gives

$$R_{i,t} - R_{f,t} = R_{f,t} \left(\exp \left(\sum_{n=2}^N \lambda_{n,t} B_{i,t}^{(n)} \right) - 1 \right) + u_{i,t}, \quad \mathbb{E}_{t-1}[u_{i,t}] = 0. \quad (\text{IA4})$$

Economic magnitudes. At monthly frequency, the own-risk correction is usually dominated by the variance term $\kappa_2(r_i)/2$, while the third- and fourth-order own-return terms are much smaller. For this reason, the paper treats Route 1 and Route 2 as supporting derivations rather than as the main organizing device of the empirical work.

IA1.2. Alpha mapping in exact and approximate form

Using $\alpha_i^{(N)} = \mathbb{K}[r_i - r_f] - \widehat{\mathbb{K}[r_i - r_f]}^{(N)}$, the exact simple-return-comparable cumulant alpha can be rewritten exactly as

$$\alpha_i^{(R,N)} = R_f e^{\widehat{\mathbb{K}[r_i - r_f]}^{(N)}} \widetilde{\alpha}_i^{(N)}, \quad \widetilde{\alpha}_i^{(N)} \equiv e^{\alpha_i^{(N)}} - 1.$$

When pricing errors are small,

$$\alpha_i^{(R,N)} \approx R_f e^{\widehat{\mathbb{K}[r_i - r_f]}^{(N)}} \alpha_i^{(N)},$$

which explains why $\widetilde{\alpha}_i^{(N)}$ is a useful diagnostic even though the main text uses the exact mapping for benchmark comparability.

IA1.3. Nonlinear factor representations, level-space nonlinearity, and affine readouts

Assumption 1 restricts the map from the selected factor representation to the negative log SDF; it does not require the factors themselves to be linear functions of primitive economic information. Let z_{t+1} denote a potentially high-dimensional information set and let

$$f_{t+1} = \Phi_{\psi,t}(z_{t+1}) \in \mathbb{R}^D \quad (\text{IA5})$$

be any $\mathcal{F}_{t+1}$ -measurable representation. The map $\Phi_{\psi,t}$ may be an economic portfolio construction, a principal-component transformation, or a nonlinear machine-learning

architecture. If the negative log SDF has the affine form

$$m_{t+1} = a_t + b_t' f_{t+1} = a_t + b_t' \Phi_{\psi,t}(z_{t+1}), \quad (\text{IA6})$$

then Propositions 2–4 apply without modification. The positive pricing kernel is

$$M_{t+1} = \exp\{-a_t - b_t' \Phi_{\psi,t}(z_{t+1})\} = e^{-a_t} \sum_{k=0}^{\infty} \frac{[-b_t' \Phi_{\psi,t}(z_{t+1})]^k}{k!}. \quad (\text{IA7})$$

Thus, a low-dimensional affine index for m_{t+1} generates a generally nonlinear pricing kernel in the primitive state. All powers and interactions in Equation (IA7) are governed by the same loading vector b_t , which is the level-space source of the monomial restrictions in Proposition 3. PCA is the special case in which $\Phi_{\psi,t}$ is linear; characteristic-based and machine-learning methods allow the representation to be nonlinear or state dependent (Kelly et al., 2019; Gu et al., 2020).

This distinction separates two uses of data-driven methods. When a method supplies a low-dimensional factor vector f^{ML} , Equation (IA6) is the maintained spanning restriction and the full monomial price structure of Proposition 3 remains testable. When a method instead supplies a direct positive SDF M^{ML} , Theorem 1 applies to $m^{\text{ML}} = -\log M^{\text{ML}}$ without an affine restriction; the monomial structure follows only if that direct SDF admits an affine bottleneck or readout.

IA1.4. Factor-coordinate equivariance

The complete affine model is invariant to a nonsingular change of factor coordinates. Let

$$f_{t+1}^* = C f_{t+1} + d, \quad C \in \mathbb{R}^{D \times D} \text{ nonsingular}, \quad (\text{IA8})$$

and define

$$b_t^* = (C')^{-1} b_t, \quad a_t^* = a_t - b_t' C^{-1} d. \quad (\text{IA9})$$

Then $m_{t+1} = a_t^* + b_t^{*'} f_{t+1}^*$. The shift d drops from every mixed cumulant of total order at least two. If $Z_{i,t}^{(p,h)}$ is the raw order- h factor-cumulant tensor and $\widetilde{Z}_{i,t}^{(p,h)}$ is its K-beta

transform, multilinearity and the covariance transformation $\Sigma_{f^\star,t} = C\Sigma_{f,t}C'$ give

$$Z_{i,t}^{\star(p,h)} = Z_{i,t}^{(p,h)} \times_1 C \times_2 \cdots \times_h C, \quad (\text{IA10})$$

$$\widetilde{Z}_{i,t}^{\star(p,h)} = \widetilde{Z}_{i,t}^{(p,h)} \times_1 (C')^{-1} \times_2 \cdots \times_h (C')^{-1}. \quad (\text{IA11})$$

Moreover,

$$\theta_t^\star \equiv \Sigma_{f^\star,t} b_t^\star = C\theta_t. \quad (\text{IA12})$$

Consequently, the complete pricing contractions are unchanged:

$$\langle (b_t^\star)^{\otimes h}, Z_{i,t}^{\star(p,h)} \rangle = \langle b_t^{\otimes h}, Z_{i,t}^{(p,h)} \rangle, \quad (\text{IA13})$$

$$\langle (\theta_t^\star)^{\otimes h}, \widetilde{Z}_{i,t}^{\star(p,h)} \rangle = \langle \theta_t^{\otimes h}, \widetilde{Z}_{i,t}^{(p,h)} \rangle. \quad (\text{IA14})$$

The individual coordinates and monomials, therefore, depend on the chosen economic factor basis, while their complete fitted contribution does not. This is the same equivariance familiar from an ordinary multivariate beta model.

IA2. Supporting Empirical Details

Standard errors for log-space alphas

Let $\widehat{\alpha}_{i,t}^{(N)}$ denote the monthly generated log-space pricing residual for asset i . The reported unconditional alpha is the time average

$$\bar{\alpha}_i^{(N)} \equiv \frac{1}{T} \sum_{t=1}^T \widehat{\alpha}_{i,t}^{(N)}.$$

Per-asset inference applies Newey–West standard errors (six lags) to the time-series average of the generated residual, as in Section 5.1.6. Because the residual is a generated quantity (both the exposures and the risk prices are estimated), inference on model comparisons in the main text relies instead on a stationary block bootstrap over months that resamples entire cross-sections and recomputes the comparison statistic inside each draw (Section 5.2.1).

Benchmark-comparable cumulant alpha construction

For benchmark comparisons, the paper also needs the generated monthly simple-return residual associated with the cumulant model. The baseline object is the exact month- t residual

$$\widehat{\alpha}_{i,t}^{(R,N)} \equiv (R_{i,t} - R_{f,t}) - R_{f,t} \left(e^{\mathbb{K}_{t-1}[\widehat{r_{i,t} - r_{f,t}}]^{(N)}} - 1 \right).$$

The reported unconditional benchmark-comparable cumulant alpha is then

$$\overline{\alpha}_i^{(R,N)} \equiv \frac{1}{T} \sum_{t=1}^T \widehat{\alpha}_{i,t}^{(R,N)}.$$

This is the object that should be compared to unconditional CAPM, FF3, and FF6 alphas in Section 5. The transformed wedge $\widetilde{\alpha}_i^{(N)}$ is retained only as a diagnostic summary, not as the benchmark comparison object.

IA3. Alternative SDF Constructions

The main text implements the affine specification of Assumption 1 for the market-only, FF3, MRC, and FF6 spans. This appendix records alternative SDF constructions that are compatible with the framework and can be carried through the same cumulant machinery.

Minimum-HJ and FF25-based constructions

One alternative uses the 25 Fama–French size-value portfolios, or a related benchmark set, to estimate an SDF by minimum-HJ or related distance criteria. The resulting kernel satisfies the Euler equation on the benchmark set by construction, so Theorem 1 applies to it directly; we leave its implementation to future work.

Exponential-tilt SDFs

A second alternative models the SDF as a low-dimensional exponential tilt on a dictionary of traded log-return features. Such constructions are economically interesting because

they allow the data to select a small set of effective risk factors, and they slot into the framework through the affine readout of Section [IA1.3](#).

Intermediary-motivated SDFs

Intermediary asset pricing provides a structural rationale for why higher-order co-cumulants may matter more when intermediary risk-bearing capacity is low: tighter balance-sheet capacity amplifies priced risks ([He and Krishnamurthy, 2013](#); [Adrian et al., 2014](#)). That logic is a useful motivation for the higher-order channels, but it is not necessary for the paper’s theorem or empirical construction, so we record it here rather than in the main text.

Machine-learning SDFs

We also consider a machine-learning construction in which a potentially high-dimensional information set is mapped into a low-dimensional representation ([Kelly et al., 2019](#); [Gu et al., 2020](#); [Kozak et al., 2020](#)). Propositions [2–4](#) thus apply without modification in the learned coordinates, and the higher-order price coefficients retain the monomial structure implied by the affine structure.

IA4. Additional Asset-Pricing Tests and Robustness

Residual-share diagnostics

A useful set of supporting diagnostics reports how much of the absolute pricing wedge is removed by each additional order:

$$1 - \frac{\text{MAE}(\alpha^{(2)})}{\text{MAE}(\mathbb{K}[r_i - r_f])}, \quad 1 - \frac{\text{MAE}(\alpha^{(3)})}{\text{MAE}(\alpha^{(2)})}, \quad 1 - \frac{\text{MAE}(\alpha^{(4)})}{\text{MAE}(\alpha^{(3)})}.$$

These ratios summarize how much of the pricing wedge each additional order removes and support the stopping rule of Section [3.5](#).

Subperiod, crisis-exclusion, and influence diagnostics

We evaluate the original past-only forecasts over three calendar subperiods: 1978–1999, 2000–2009, and 2010–2024. We also remove the 2008–2009 financial crisis window and the March–June 2020 market dislocation window from the full evaluation panel. In each case, the forecast series remains fixed. Only the dates entering the evaluation statistic change. Per-asset HAC statistics are recomputed within the stated evaluation sample.

In the untabulated table, $\mathbb{K}$ -Beta[6, 4] yields a lower mean absolute alpha, a lower cross-sectional mean of squared HAC alpha statistics, and a lower evaluated Fama–MacBeth intercept than raw FF6 in each of the three subperiods and in both crisis-exclusion samples. The magnitude of the mean-absolute-alpha improvement varies across evaluation windows. The improvements in alpha-significance and intercept are directionally uniform. Excluding 2008–2009 or March–June 2020 preserves the full-sample ranking on all three diagnostics.

Higher-order truncation checks

The fourth-order specification leaves the truncation residual formed by cumulant orders five and above. At each additional total order $P \in \{5, 6\}$, we construct every transformed mixed co-cumulant with $p + |q| = P$, apply the Euler-equation coefficient $c_{p,q}$, aggregate the terms into one order block, and estimate one additional monthly Fama–MacBeth price. The recursive moment half-lives, minimum-history rules, factor-covariance normalization, and trailing sixty-month price forecast remain the same as in the fourth-order specification.

In untabulated diagnostics, their estimated contributions are economically negligible relative to $\mathbb{K}[r_i - r_f]$ and to the order-four residual, consistent with the factorial decay of the combinatorial weights $c_{a,b}$; this supports the fourth order as the stopping point used in the main text (Section 3.5).

IA5. Notation

Terminology

Term	Meaning in this paper
$\mathbb{K}$ -expected excess log return	$\mathbb{K}[r_i - r_f]$, not the arithmetic mean of $r_i - r_f$.
Exact-affine baseline SDF	The rolling affine-in-market SDF $m_t = a_t + b_t r_{m,t}$ identified from the market and risk-free $\mathbb{K}$ -expectation restrictions.
Mixed co-cumulant	A co-cumulant $\kappa_{a,b}(r_i, m)$ with $a, b \geq 1$; the dependence terms (the latent-SDF cumulant factors).
Risk factor	One of the D observable variables f_d that span the SDF (Assumption 1); few in number.
Cumulant factor	A joint-cumulant term $Q_{i,p,q} = c_{p,q} \kappa_{p,q}(r_i, f)$ of returns with the risk factors; the building blocks the expansion generates. Combinatorially many.
Cumulant-factor library	The full set of $J(D, P)$ cumulant factors at truncation order P ; the theoretical object, distinct from what is estimated.
Priced cumulant factor	A cumulant factor that earns a nonzero premium in the cross section; “priced” is reserved for this empirical sense, not for carrying a structural price coefficient.
Structural price coefficient	$\eta_{p,q} = \prod_d b_d^{q_d}$ (or θ^q in $\mathbb{K}$ -beta coordinates); nonzero under the affine null when all $b_d \neq 0$, distinct from an earned premium.
Pseudo-price of co-cumulant risk	$\pi_{a,b} \equiv c_{a,b} \kappa_{a+b}(m)$; theory-implied scaling, not a dollar price.
Order-level price of risk	$\lambda_n \equiv \kappa_n(m)$ for total order n .
Composite beta	$B_i^{(n)}$, the within-order weighted sum of normalized co-cumulant betas.
Own-risk correction	Ω_i , the return-only cumulant correction linking $\mathbb{K}[r_i - r_f]$ to $\mathbb{E}[r_i - r_f]$.
Log-space alpha	A pricing error measured in log units, reported in lbp.
Multiplicative factor structure	The linear-space representation $\mathbb{E}[R_i] = R_f \prod_n \Lambda_n^{B_i^{(n)}}$.
Canonical portfolio direction	The normalized convention in which P1 is the economically short endpoint and P5 is the economically long endpoint for every signal.
lbp	Log-base-points; one lbp is 10^{-4} log units.

Time and expectations

Symbol	Meaning
t	Time index.
T	Number of time observations.
$\mathcal{F}_t$	Information set at time t .
$\mathbb{E}_t[\cdot]$	Conditional expectation operator $\mathbb{E}[\cdot \mid \mathcal{F}_t]$.
$\mathbb{K}_t[\cdot]$	Conditional $\mathbb{K}$ -expectation $\log \mathbb{E}[\exp(\cdot) \mid \mathcal{F}_t]$.

Returns and SDF

Symbol	Meaning
$M_t > 0$	Stochastic discount factor.
$m_t \equiv -\log M_t$	Log-SDF.
$R_{i,t} > 0$	Gross return on asset i .
$R_{m,t} > 0$	Gross market return.
$R_{f,t} > 0$	Gross risk-free return, known at $t - 1$.
$r_{i,t} \equiv \log R_{i,t}$	Log return on asset i .
$r_{m,t} \equiv \log R_{m,t}$	Log market return.
$r_{f,t} \equiv \log R_{f,t}$	Log risk-free return.
$\mathbb{K}_{t-1}[r_{i,t} - r_{f,t}]$	$\mathbb{K}$ -expected excess log return.
a_t, b_t	Rolling exact-affine baseline SDF coefficients in $m_t = a_t + b_t r_{m,t}$.

Cumulants and pricing objects

Symbol	Meaning
$\kappa_n(X)$	n th cumulant of X .
$\kappa_{a,b}(X, Y)$	Co-cumulant of order (a, b) for (X, Y) .
$c_{a,b}$	Theory-implied coefficient $(-1)^{b+1}/(a! b!)$.
$\beta_i^{(a,b)}$	Normalized co-cumulant beta $\kappa_{a,b}(r_i, m)/\kappa_{a+b}(m)$.
$B_i^{(n)}$	Composite beta at total order n .
λ_n	Order-level price of risk $\kappa_n(m)$.
$\pi_{a,b}$	Pseudo-price of co-cumulant risk, equal to $c_{a,b}\kappa_{a+b}(m)$.
Λ_n	Linear-space factor $\exp(\lambda_n)$.
$f = (f_1, \dots, f_D)$	Observable risk factors spanning the SDF.
b_d	SDF loading on risk factor f_d in $m = a + b'f$.
$Q_{i,p,q}, \tilde{Q}_{i,p,q}$	Cumulant factor $c_{p,q}\kappa_{p,q}(r_i, f)$ and its $\mathbb{K}$ -beta version.
$\eta_{p,q}$	Structural price coefficient $\prod_d b_d^{q_d}$ (θ^q in $\mathbb{K}$ -beta coordinates).
$J(D, P)$	Size of the cumulant-factor library at order P , $\sum_{h=1}^{P-1} (P-h) \binom{D+h-1}{h}$.
D_N	Dimension of the mixed co-cumulant space through order N , equal to $\binom{N}{2}$.
$\mathbb{K}[r_i - r_f]^{(N)}$	Truncated order- N model-implied value of $\mathbb{K}[r_i - r_f]$.
$\alpha_i^{(N)}$	Log-space pricing error $\mathbb{K}[r_i - r_f] - \mathbb{K}[r_i - r_f]^{(N)}$.
$\zeta_i^{(N)}$	Truncation residual of the mixed co-cumulant expansion.
Ω_i	Own-risk correction linking $\mathbb{K}[r_i - r_f]$ to $\mathbb{E}[r_i - r_f]$.
$\tilde{\alpha}_{i,t}^{(R,N)}$	Generated monthly simple-return residual associated with the order- N cumulant model.
$\tilde{\alpha}_i^{(R,N)}$	Unconditional mean of the generated simple-return residual.
$\alpha_i^{(R,N)}$	Exact simple-return-comparable cumulant alpha.
$\tilde{\alpha}_i^{(N)}$	Exponential wedge $\exp(\alpha_i^{(N)}) - 1$, prior to the exact $R_f e^{\widehat{\mathbb{K}[r_i - r_f]}^{(N)}}$ scaling.
$\lambda_5, \lambda_6, B_i^{(5)}, B_i^{(6)}$	Diagnostic higher-order objects used to assess whether the order-four stopping rule is economically adequate.

IA6. Anomalies and Family Mapping

Family	Anomalies: Notation as in Chen and Zimmermann (2022)
Risk/vol	BetaBDLeverage, betaCC, betaCR, BetaDimson, BetaFP, betaNet, betaRC, betaRR, BetaSquared, BetaTailRisk, betaVIX, CoskewACX, Coskewness, DownsideBeta, IdioVol3F, IdioVolAHT, IdioVolCAPM, IdioVolQF, MaxRet, MomVol, RealizedVol, ReturnSkew, ReturnSkew3F, ReturnSkewCAPM, ReturnSkewQF, roavol, skew1, VarCF, VolMkt, VolSD
Size/age	AgeIPO, FirmAge, Size
Value	AOP, BM, BMdec, BMq, BPEBM, cfp, cfpq, CFq, EBM, EBM_q, EP, EPq, EquityDuration, IntrinsicValue, Price, SP, SP_q
Profitability	cashdebt, CashProd, CBOperProf, CBOperProfLagAT, CBOperProfLagAT_q, ChangeRoA, ChangeRoE, GP, GPlag, GPlag_q, LaborforceEfficiency, OperProf, OperProfLag, OperProfLag_q, OperProfRD, OperProfRDLagAT, OperProfRDLagAT_q, OPLeverage, OPLeverage_q, PM_q, roaq, RoE, roic, sfe
Investment	AssetGrowth, AssetGrowth_q, AssetTurnover, AssetTurnover_q, BrandCapital, BrandInvest, CapTurnover, CapTurnover_q, ChAssetTurnover, ChInv, ChInvIA, depr, dNoa, grcapx, grcapx1y, grcapx3y, GrLTNOA, hire, Investment, InvestPPEInv, InvGrowth, MeanRankRevGrowth, NOA, OrgCap, OrgCapNoAdj, pchdepr, realestate, RetNOA, RetNOA_q, sgr, sgr_q
Momentum	CustomerMomentum, FirmAgeMom, High52, IndMom, IndRetBig, IntMom, iomom_cust, iomom_supp, Mom12m, Mom6m, Mom6mJunk, ResidualMomentum, ResidualMomentum6m, retConglomerate, TrendFactor
Earnings quality	AnalystRevision, AnalystValue, AnnouncementReturn, ChangeInRecommendation, ChPM, ChTax, EarningsConservatism, EarningsConsistency, EarningsForecastDisparity, EarningsPersistence, EarningsPredictability, EarningsSmoothness, EarningsSurprise, EarningsTimeliness, EarningsValueRelevance, ETR, FEPS, fgr5yrLag, fgr5yrNoLag, ForecastDispersion, ForecastDispersionLT, FR, FRbook, Frontier, PredictedFE, RDAbility, RevenueSurprise, Tax, Tax_q
Other	AM, AMq, Beta, Cash, CF, ChEQ, ChNCOA, ChNCOL, ChNNCOA, EarnSupBig, ExclExp, GrGMToGrSales, GrSaleToGrInv, GrSaleToGrOverhead, GrSaleToGrReceivables, pchgm_pchsale, pchsaleinv, PM, PS, PS_q, salecash, saleinv, salerec
Distress	BookLeverage, BookLeverageQuarterly, currat, EntMult, EntMult_q, FailureProbability, FailureProbabilityJune, KZ, KZ_q, Leverage, Leverage_q, NetDebtPrice, NetDebtPrice_q, pchcurrat, pchquick, quick, tang, tang_q, WW, WW_Q, ZScore, ZScore_q
Issuance	CompEqulss, CompositeDebtIssuance, DelCOA, DelCOL, DelDRC, DelEqu, DelFINL, DelLTI, DelNetFin, DelSTI, NetDebtFinance, NetEquityFinance, NetPayoutYield, NetPayoutYield_q, PayoutYield, PayoutYield_q, ShareIss1Y, ShareIss5Y, XFIN
Liquidity	AssetLiquidityBook, AssetLiquidityBookQuart, AssetLiquidityMarket, AssetLiquidityMarketQuart, BetaLiquidityPS, BidAskSpread, DelayAcct, DelayNonAcct, DolVol, Illiquidity, PriceDelayRsqr, PriceDelaySlope, PriceDelayTstat, std_turn, VolumeTrend, zero-trade12M, zero-trade1M, zero-trade6M
Intangibles	AdExp, GrAdExp, Herf, HerfAsset, HerfBE, IntanBM, IntanCFP, IntanEP, IntanSP, OrderBacklog, OrderBacklogChg, RD, rd_sale, rd_sale_q
Seasonality	Mom12mOffSeason, MomOffSeason, MomOffSeason06YrPlus, MomOffSeason11YrPlus, MomOffSeason16YrPlus, MomSeason, MomSeason06YrPlus, MomSeason11YrPlus, MomSeason16YrPlus, MomSeasonShort
Accruals	AbnormalAccruals, AbnormalAccrualsPercent, AccrualQuality, AccrualQualityJune, Accruals, ChNWC, PctAcc, PctTotAcc, TotalAccruals
Options vol.	CPVolSpread, dCPVolSpread, dVolCall, dVolPut, OptionVolume1, OptionVolume2, ProbInformedTrading, RIVolSpread, SmileSlope
Reversal	EarningsStreak, LRreversal, MRreversal, REV6, STreversal
Short interest	DelBreadth, IO_ShortInterest, RIO_Disb, RIO_Turnover, ShortInterest