When Does a High Sharpe Ratio Imply a Risk Premium?

March 19, 2026

Abstract

A factor can earn a large expected return while its beta carries zero or negative cross-sectional price, because earning a return and pricing the cross-section are logically distinct. Projecting the SDF onto any portfolio decomposes the Fama-MacBeth slope into two components: the factor's expected return, which the Sharpe ratio governs, and a cross-sectional alignment term, the covariance of pricing errors with betas, which the Sharpe ratio does not govern. The alignment term further decomposes into a magnitude the GRS alpha budget constrains and a directional component, the cross-sectional correlation ρ between betas and pricing errors, that is unconstrained. A factor can have an arbitrarily high Sharpe ratio while its Fama-MacBeth slope takes any value, including the wrong sign. The same alignment term shifts the intercept in the opposite direction through the centroid constraint: the flat security market line and the elevated zero-beta rate are one finding. Across 212 factors and 1,277 test portfolios, the correlation between absolute Sharpe ratios and Fama-MacBeth slopes never exceeds 0.19; wrong-sign slopes are common even among high-Sharpe-ratio factors. The Sharpe ratio ranks models by aggregate fit but does not determine cross-sectional prices of risk.

1 Introduction

The market earns roughly 8% per year, yet market beta carries almost no cross-sectional price. The Fama-MacBeth slope on market beta is close to zero even when betas are estimated precisely from daily returns. These are different facts about different objects: the market's expected return is a time-series property of one portfolio; the Fama-MacBeth slope is a cross-sectional property of how expected returns vary with beta across many assets. The literature treats them as the same object and interprets the discrepancy as a puzzle. This paper shows they are generically different, and explains what determines each. The regression intercept, which exceeds the risk-free rate by a wide margin, is a corollary of the same decomposition: the regression line passes through the centroid $(\bar{\beta}, \bar{\mu})$, so the same force that compresses the slope inflates the intercept. The flat security market line and the elevated zero-beta rate are one finding, not two (Black 1972; Frazzini and Pedersen 2014).

The anatomy of the Fama-MacBeth slope makes the distinction precise. Project the stochastic discount factor m onto the market return: $m = a + b\, r_m + \varepsilon$, where ε is the SDF component orthogonal to the market. Substituting into the pricing condition $E[m \cdot r_i] = 0$ gives $E[r_i] = \beta_{i,m}\, E[r_m] + \alpha_{i,m}$, where $\alpha_{i,m} = -\operatorname{Cov}(\varepsilon, r_i)/E[m]$ captures each asset's exposure to the unspanned SDF. The Fama-MacBeth regression cannot control for $\alpha_{i,m}$ because ε is unobservable, so γ_1 has two components: the factor's expected return $E[r_m]$, which the Sharpe ratio governs, and the cross-sectional covariance of pricing errors with betas, which the Sharpe ratio does not govern. The second component is a population quantity, not an estimation artifact: it persists with infinite samples, precise betas, and GLS weighting (equation (11)). The intercept absorbs the mirror image through the centroid constraint (equation (28)). The flat security market line requires only that the market does not span the SDF, which is the maintained assumption of every multi-factor model.

The result is not specific to the market. Project the SDF onto any portfolio p with a nonzero expected return: the same decomposition applies, with ε representing the SDF component unspanned by p. Every portfolio with $E[r_p] \neq 0$ has a nonzero projection coefficient

$b = -E[m]\, E[r_p]/\,\mathrm{Var}(r_p)$, so every such portfolio is "in the SDF" in the sense of being correlated with the pricing kernel, but this bar is trivially low: it says nothing about whether the portfolio's beta carries a cross-sectional price. The Fama-MacBeth slope γ_1 always has two components: $E[r_p]$ and a cross-sectional alignment term that depends on how the unspanned SDF component covaries with each test asset across the beta distribution. The maximum squared Sharpe ratio has become the primary criterion for comparing asset pricing models.[1] A higher Sharpe ratio means the factor captures more SDF variance, tightening the aggregate mispricing budget. But the alignment term depends not on the aggregate budget but on how mispricing distributes across the cross-section. Corollary 1 decomposes the alignment term into $\rho_{\mathrm{cs}}(\beta, \alpha) \cdot \mathrm{SD}_{\mathrm{cs}}(\alpha)/\mathrm{SD}_{\mathrm{cs}}(\beta)$. The Sharpe ratio constrains the magnitude $\mathrm{SD}_{\mathrm{cs}}(\alpha)$ through the GRS alpha budget, but the directional component $\rho_{\mathrm{cs}}(\beta, \alpha)$, the cross-sectional correlation between betas and pricing errors, is unconstrained. Whether a given Sharpe ratio qualifies as "high" depends on the unobservable tangency Sharpe ratio (Hansen and Jagannathan 1991; Hansen and Richard 1987); a Cauchy-Schwarz bound links the two (Proposition 1), but the bound is generically loose. SDF spanning is the only condition that guarantees the two components coincide for all test assets (Corollary 3).

The result extends to multi-factor models through the Frisch-Waugh-Lovell (FWL) theorem. In a K-factor model, the OFB for factor k depends not on raw cross-sectional variation in $\beta_{i,k}$ but on the partialled-out component: the variation unexplained by betas on all other factors (Proposition 2). Correlated factors produce correlated betas across assets, reducing the FWL denominator and amplifying the alignment term per unit of mispricing. Adding factors improves the model's aggregate Sharpe ratio by expanding the spanned subspace but simultaneously increases multicollinearity among betas, so the net effect on the OFB for any individual factor is ambiguous.

1. (Kozak, Nagel, and Santosh 2018, 2020; Barillas et al. 2020; Fama and French 2018; Lettau and Pelger 2020; Chen, Pelger, and Zhu 2024) Barillas and Shanken (2017) show that "the relative pricing performance of two models can be evaluated simply by comparing their squared Sharpe ratios." Barillas et al. (2020) formalize this through a squared Sharpe ratio improvement measure, and Fama and French (2018) adopt the same metric.

Section 4 uses 212 long-short factor portfolios from the Chen and Zimmermann (2022) Open Source Asset Pricing database, treating each as a candidate single-factor model. The test assets are 1,277 sorted portfolios from the same database. Using four beta estimation methods (full-sample OLS, rolling 60-month, daily 252-day, and daily 1260-day windows), the correlation between absolute Sharpe ratios and Fama-MacBeth slopes never exceeds 0.19. Between 22% and 31% of factors produce wrong-sign slopes. The wrong-sign rate drops from 44% among the lowest Sharpe-ratio quintile to 2% among the highest, but the alignment term does not shrink: the Sharpe ratio predicts the sign of γ_1 but not its magnitude. Full-sample cross-sectional slopes, which eliminate estimation error by construction, confirm that the alignment term rather than measurement noise drives the disconnect between what factors earn and what their betas price. Multi-factor models exhibit the same pattern, with FWL amplification producing large alignment terms for individual factors even when the model's aggregate Sharpe ratio is high.

1.1 Related literature

The omitted variable bias formula applied to a cross-sectional regression is standard (Cochrane 2005). The contribution is to derive it through the SDF projection, which connects the alignment term directly to the factor's Sharpe ratio through the GRS budget and to spanning through $\varepsilon = 0$. This connection is what makes the Sharpe ratio question precise: the SR constrains $\mathrm{Var}(\varepsilon)$ and hence the magnitude of mispricing, but the alignment term depends on how that mispricing distributes across the cross-section of betas, a directional property the SR does not govern. The decomposition of the alignment term into magnitude ($\mathrm{SD}_{\mathrm{cs}}(\alpha)/\mathrm{SD}_{\mathrm{cs}}(\beta)$) and direction ($\rho_{\mathrm{cs}}(\beta, \alpha)$) in Corollary 1, the calibration showing the Cauchy-Schwarz bound is loose for realistic parameters (Proposition 1), the FWL amplification in multi-factor models (Proposition 2), and the systematic empirical documentation across 212 factors are, to my knowledge, new. The conceptual point that a factor's expected return and its cross-sectional price are distinct objects, connected only under spanning, reframes the flat security market line as a structural feature of incomplete models rather than a puzzle requiring a behavioral explanation.

Giglio and Xiu (2021) propose a three-pass procedure to estimate risk premia when priced factors are omitted, assuming that the factors driving return covariation are pervasive enough for PCA to recover them. Their approach makes the two components of γ_1 coincide for factors well-spanned by PCA; the decomposition characterizes why they differ in general. Giglio, Xiu, and Zhang (2021) extend the three-pass estimator to weak factors. The present paper asks a different question: not how to make the two components coincide, but what determines the alignment term's sign and magnitude, and why the Sharpe ratio fails to discipline it. Bryzgalova et al. (2023) define tradable factor risk premia through the SDF projection onto the full return space rather than the factor space, obtaining risk premia that are invariant to model composition; for tradable factors, these reduce to the factor's expected return, sidestepping the model-dependence the FWL decomposition (Proposition 2) predicts for conventional slopes.

The broader literature has addressed related but distinct problems. Kandel and Stambaugh (1995) show through mean-variance geometry that OLS cross-sectional slopes depend on the test-asset rotation when the factor is inefficient and propose GLS as a rotation-invariant alternative. The present paper differs in approach and conclusion: the SDF projection derives the alignment term algebraically (equation (11)) and shows that GLS inherits it, because the source is misspecification, not the estimator's weighting scheme. Rotation invariance does not eliminate the alignment term; it merely ensures the same nonzero value across affine transformations of the test assets. Shanken (1987) bounds total mispricing when a proxy replaces the true market portfolio. Kan and Zhang (1999) show that useless factors produce spuriously significant t-statistics. Kan, Robotti, and Shanken (2013) and Gospodinov, Kan, and Robotti (2014) develop misspecification-robust inference. Lewellen, Nagel, and Shanken (2010) argue that narrow test assets produce misleading diagnostics. Kozak and Nagel (2024) characterize when heuristic factor portfolios span the SDF. Giannetti (2024) proposes a Hausman test detecting whether the Fama-MacBeth slope equals the factor's expected return. Di Tella et al. (2026) construct a time-varying zero-beta rate from the Fama-MacBeth intercept and show it fits the aggregate consumption Euler equation; the decomposition (equation (28)) suggests that without spanning, the intercept conflates any true zero-beta premium with the mechanical seesaw from the alignment term, so the spanning assumption their identification requires should be verified empirically. Wang (2024) shows that the minimum-variance zero-beta portfolio converges toward the global minimum-variance portfolio under misspecification, providing a portfolio-geometry explanation for the elevated intercept; the seesaw identity complements this result by tracing the same phenomenon to the cross-sectional correlation between betas and pricing errors, and by showing the intercept and slope deviations are one bias, not two. These papers address inference, specification testing, or correction; none decomposes the alignment term into a magnitude component the Sharpe ratio constrains and a directional component it does not, or quantifies the disconnect empirically across the factor zoo.

The Sharpe-ratio paradigm for model comparison is the most immediate context. The theorems of Barillas and Shanken (2017), Barillas et al. (2020), Kozak, Nagel, and Santosh (2018), and Fama and French (2018) are correct: a higher SR^2 does mean smaller aggregate alphas, and this paper does not dispute those results. But the language of "pricing performance" invites the further inference that the winning model's Fama-MacBeth slopes coincide with the factors' expected returns, and that inference does not follow. The formal results concern aggregate fit; the gap between aggregate fit and individual cross-sectional prices is the subject of this paper. The Sharpe ratio governs the alpha vector's norm, not its direction. Harvey, Liu, and Zhu (2016) and McLean and Pontiff (2016) document that many published anomalies reflect data mining or attenuate out of sample; the OFB results here do not depend on whether the 212 factors represent genuine risk or spurious patterns, because the identity holds for any portfolio with a nonzero expected return. Easterwood and Paye (2025) show that estimated maximum Sharpe ratios for prominent models are upward biased by data snooping. Chernov, Dahlquist, and Lochstoer (2025) document that unpriced components account for 30 to 99 percent of factor return variance; the cross-sectional distribution of this unpriced risk governs the OFB. Hollstein and Prokopczuk (2022) confront seven factor models with the Fama-MacBeth test across a wide range of test assets and beta estimators, finding risk premium estimates well below average factor returns, patterns the OFB formula predicts. The foundational results on SDF bounds and projections are due to Hansen and Jagannathan (1991) and Hansen and Richard (1987); Cochrane (2005) provides the textbook treatment.

2 Theory

2.1 Setup

Consider an economy with N assets having excess returns r_i, $i = 1, \ldots, N$, and a stochastic discount factor m satisfying $E[m \cdot r_i] = 0$ for all i. Let r_p denote the excess return on an arbitrary portfolio p, and define its Sharpe ratio as $\mathrm{SR}(p) = E[r_p]/\mathrm{SD}(r_p)$.

2.2 SDF projection onto an arbitrary portfolio

Project m onto the space spanned by r_p:

$$m = a + b\, r_p + \varepsilon, \qquad \text{where } \operatorname{Cov}(\varepsilon, r_p) = 0. \tag{1}$$

The projection coefficient is

$$b = \frac{\operatorname{Cov}(m, r_p)}{\operatorname{Var}(r_p)}. \tag{2}$$

From the pricing condition $E[m \cdot r_p] = 0$:

$$E[m]\, E[r_p] + \operatorname{Cov}(m, r_p) = 0 \quad \Longrightarrow \quad b = -\frac{E[m]\, E[r_p]}{\operatorname{Var}(r_p)}.$$

2.3 Pricing with an incomplete factor model

Substitute the projection (1) into the pricing condition for any excess return r_i:

$$\begin{aligned} 0 &= E[m \cdot r_i] = E\big[(a + b\, r_p + \varepsilon)\, r_i\big] \\ &= E[m]\, E[r_i] + b\ \mathrm{Cov}(r_p, r_i) + \mathrm{Cov}(\varepsilon, r_i), \end{aligned} \tag{3}$$

where I used $E[m] = a + b\, E[r_p]$ (from taking expectations of (1) with $E[\varepsilon] = 0$). Dividing by $-E[m]$ and defining $\beta_{i,p} = \mathrm{Cov}(r_p, r_i)/\operatorname{Var}(r_p)$:

$$E[r_i] = \beta_{i,p}\, E[r_p] - \frac{\mathrm{Cov}(\varepsilon, r_i)}{E[m]}. \tag{4}$$

The coefficient on $\beta_{i,p}$ is $E[r_p]$, the factor's own expected return. But the equation also contains the term $\mathrm{Cov}(\varepsilon, r_i)/E[m]$, which captures the pricing contribution of every SDF component not spanned by r_p. The coefficient $E[r_p]$ is the cross-sectional price of $\beta_{i,p}$ only when this second term is absent, that is, only when r_p spans the SDF.

2.4 Anatomy of the Fama-MacBeth slope

In a univariate cross-sectional regression of $E[r_i]$ on $\beta_{i,p}$ (with an intercept), the OLS slope is $\gamma_1 = \mathrm{Cov}_{\mathrm{cs}}(\beta_{i,p},\ E[r_i])/\operatorname{Var}_{\mathrm{cs}}(\beta_{i,p})$, where $\mathrm{Cov}_{\mathrm{cs}}$ and $\mathrm{Var}_{\mathrm{cs}}$ denote cross-sectional moments over $i = 1, \ldots, N$. Substituting the pricing equation (4):

$$\gamma_1 = E[r_p] - \frac{\mathrm{Cov}_{\mathrm{cs}}(\beta_{i,p},\ \mathrm{Cov}(\varepsilon, r_i)/E[m])}{\operatorname{Var}_{\mathrm{cs}}(\beta_{i,p})}, \tag{5}$$

since the $\beta_{i,p}\, E[r_p]$ term in (4) contributes $E[r_p]\ \operatorname{Var}_{\mathrm{cs}}(\beta_{i,p})$ to the numerator.

The Fama-MacBeth slope has two components. The first is $E[r_p]$, the factor's expected return, a time-series property of the factor portfolio alone. The second is the *omitted factor bias* (OFB), or alignment term: the cross-sectional regression coefficient of each asset's unspanned SDF exposure on its factor beta. The OFB depends on how the residual ε covaries with each asset r_i, and how those covariances line up with the betas $\beta_{i,p}$ across assets. This component can be positive, negative, or zero, regardless of the properties of p. It captures the cross-sectional interaction between the factor, the unspanned SDF component, and the test assets. The Sharpe ratio governs the first component; the second is outside its reach.

Remark 1 (Notation)**.** *The asset pricing literature uses a single symbol* λ *and a single term "risk premium" for two objects: the factor's expected return* $E[r_p]$*, a time-series property of the factor alone, and the Fama-MacBeth cross-sectional slope* γ_1*, a joint property of the factor, the test assets, and the model (Cochrane 2005). Under correct specification these coincide, so no separate terminology was needed. Under misspecification they diverge. This paper uses* γ_1 *for the FM slope and reserves* $E[r_p]$ *for the factor mean. They measure different things:* $E[r_p]$ *is what an investor earns from holding the factor portfolio;* γ_1 *is the cross-sectional price of exposure to that factor. The decomposition* (5) *shows that* γ_1 *depends on both the factor's return and a cross-sectional alignment term that the factor's own properties do not determine.*

2.5 The Fama-MacBeth slope as a portfolio return

The Fama-MacBeth slope has a direct portfolio interpretation that clarifies the economic content of the OFB. At each date t, the second-pass cross-sectional regression

$$r_{it} = \gamma_{0t} + \gamma_{1t}\,\hat{\beta}_i + e_{it}$$

produces a slope γ_{1t} via the OLS formula:

$$\gamma_{1t} = \sum_{i=1}^{N} w_i\, r_{it}, \qquad w_i = \frac{\tilde{\beta}_i}{\sum_{j=1}^{N} \tilde{\beta}_j^2}, \tag{6}$$

where $\tilde{\beta}_i = \hat{\beta}_i - \bar{\beta}$ is the demeaned beta. The time series $\{\gamma_{1t}\}$ is the return on a portfolio with weights w.

This portfolio has two properties. First, the weights sum to zero: $\sum_i w_i = \sum_i \tilde{\beta}_i / \sum_j \tilde{\beta}_j^2 = 0$, since demeaned betas sum to zero by construction. The Fama-MacBeth slope is therefore a zero-cost portfolio, going long assets with above-average betas and short those with below-average betas. Second, the portfolio has unit beta on the factor:

$$\sum_{i=1}^{N} w_i \, \beta_i = \frac{\sum_i \tilde{\beta}_i \, \beta_i}{\sum_j \tilde{\beta}_j^2} = \frac{\sum_i \tilde{\beta}_i^2}{\sum_j \tilde{\beta}_j^2} = 1, \tag{7}$$

where the second equality uses $\sum_i \tilde{\beta}_i \, \beta_i = \sum_i \tilde{\beta}_i \, (\tilde{\beta}_i + \bar{\beta}) = \sum_i \tilde{\beta}_i^2$.

The Fama-MacBeth slope is therefore the return on a zero-cost, unit-beta mimicking portfolio constructed from the test assets. Its expected return is

$$E[\gamma_{1t}] = \sum_i w_i \, E[r_i] = \sum_i w_i \, (\beta_i \, E[r_p] + \alpha_i) = E[r_p] + \sum_i w_i \, \alpha_i, \tag{8}$$

using (7) and the pricing equation $E[r_i] = \beta_i \, E[r_p] + \alpha_i$. The OFB is the alpha of the mimicking portfolio:

$$\text{OFB} = \sum_i w_i \, \alpha_i = \frac{\sum_i \tilde{\beta}_i \, \alpha_i}{\sum_i \tilde{\beta}_i^2} = \frac{\text{Cov}_{\text{cs}}(\beta, \alpha)}{\text{Var}_{\text{cs}}(\beta)}. \tag{9}$$

This expression recovers equation (5) but with a different interpretation. The mimicking portfolio replicates unit exposure to the factor but, because it is constructed from test assets that carry pricing errors, it acquires unintended exposure to the unspanned SDF component ε. The OFB is the expected return compensation for this unintended exposure.

Multi-factor case. In a K-factor cross-sectional regression, the Frisch-Waugh-Lovell theorem implies that the slope on factor k equals the return on a portfolio with weights

$$w_i^{(k)} = \frac{\tilde{\beta}_{i,k}^{\perp}}{\sum_j (\tilde{\beta}_{j,k}^{\perp})^2}, \tag{10}$$

where $\tilde{\beta}^{\perp}_{i,k}$ is the residual from regressing $\tilde{\beta}_{i,k}$ on all other demeaned betas. This portfolio has unit beta on factor k and zero beta on every other included factor: it is a "pure play" on factor k. By the same argument as (8), its alpha equals the OFB for factor k in equation (41). The FWL amplification documented in Section 2.13 now has a portfolio interpretation: when betas are correlated across factors, purging the other factors concentrates the mimicking portfolio into a smaller set of assets, increasing its exposure to any pricing errors those assets carry.

Because the OFB in (5) depends on $\mathrm{Cov}_{\mathrm{cs}}(\beta_{i,p}, \mathrm{Cov}(\varepsilon, r_i))$, a cross-sectional moment over the assets used in the regression, the Fama-MacBeth slope is not solely a property of the factor p.[2]

The OFB does not depend on the estimator's weighting scheme. The population GLS estimator (with intercept) satisfies

$$\gamma_1^{\mathrm{GLS}} = E[r_p] + \frac{\tilde{\beta}'\Sigma_e^{-1}\alpha}{\tilde{\beta}'\Sigma_e^{-1}\tilde{\beta}}, \tag{11}$$

where $\tilde{\beta}_i = \beta_{i,p} - \bar{\beta}_p$ is the demeaned beta, Σ_e is the residual covariance matrix from the time-series regressions (12), and $\alpha_i = E[r_i] - \beta_{i,p}\, E[r_p]$ is the pricing error. The expression follows from substituting $E[r_i] = \beta_{i,p}\, E[r_p] + \alpha_i$ into the GLS formula. The alignment term vanishes if and only if $\alpha = 0$ (exact pricing) or if the pricing errors are orthogonal to the betas in the Σ_e^{-1} metric. The source is not the estimator's weighting scheme but the misspecification itself: the pricing equation (4) contains the term $\mathrm{Cov}(\varepsilon, r_i)/E[m]$, and no cross-sectional estimator can remove a term that belongs in the population regression.

2.6 What determines the alignment term

The OFB formula (5) contains the term $\mathrm{Cov}(\varepsilon, r_i)$ for each test asset i. To understand what drives this term, decompose each test asset return using the time-series projection onto the

2. Kandel and Stambaugh (1995) establish a related result through mean-variance geometry.

factor:

$$r_i = \alpha_i + \beta_{i,p}\, r_p + e_i, \qquad \text{Cov}(e_i, r_p) = 0, \tag{12}$$

where $\alpha_i = E[r_i] - \beta_{i,p}\, E[r_p]$ is the time-series intercept and e_i is the residual. Since $\text{Cov}(\varepsilon, r_p) = 0$ by construction of the SDF projection (1):

$$\text{Cov}(\varepsilon, r_i) = \text{Cov}(\varepsilon, e_i). \tag{13}$$

The covariance between the SDF residual and a test asset depends entirely on the asset's residual return, the component unexplained by the factor. Substituting into (5):

$$\gamma_1 = E[r_p] - \frac{\text{Cov}_{\text{cs}}(\beta_{i,p},\ \text{Cov}(\varepsilon, e_i)/E[m])}{\text{Var}_{\text{cs}}(\beta_{i,p})}. \tag{14}$$

This decomposition clarifies the mechanism through which test assets affect the Fama-MacBeth slope. The alignment term is driven by how the residual returns e_i, the components of the test assets that the factor does not explain, covary with the SDF residual ε, the components of the SDF that the factor does not capture. If this covariance is the same for all test assets regardless of their beta, the alignment term is zero and $\gamma_1 = E[r_p]$. If it varies systematically with beta, the alignment term can be large.

To make the dependence on omitted factors explicit, suppose the SDF residual reflects a single omitted factor f_2:

$$\varepsilon = b_2\, f_2 + \eta, \qquad \text{Cov}(\eta, r_p) = \text{Cov}(\eta, f_2) = 0.$$

Then $\text{Cov}(\varepsilon, e_i) = b_2\, \text{Cov}(f_2, e_i) + \text{Cov}(\eta, e_i)$. If the test assets have no exposure to the deeper residual η, the OFB reduces to

$$\text{OFB} = \frac{b_2}{E[m]} \cdot \frac{\text{Cov}_{\text{cs}}(\beta_{i,p},\ \text{Cov}(f_2, e_i))}{\text{Var}_{\text{cs}}(\beta_{i,p})}. \tag{15}$$

The first term, $b_2/E[m]$, measures the omitted factor's importance in the SDF. The second term is the cross-sectional covariance between factor betas and the residual loadings on the omitted factor. The first is a property of the economy; the second depends entirely on which assets the researcher includes in the cross-section.

Three cases illustrate the range of outcomes.

When the test assets are sorted on the factor itself (e.g., decile portfolios formed on the characteristic underlying r_p), the extreme portfolios have mechanically high and low betas because they overlap with the long and short legs of the factor. In the data, intermediate portfolios nearly always have intermediate betas: across the 212 signals, the rank correlation between portfolio number and own-factor beta exceeds 0.8 for 95% of factors. If the omitted factor f_2 loads on the sorting characteristic in a way that is balanced across beta groups, $\mathrm{Cov}_{\mathrm{cs}}(\beta_{i,p}, \mathrm{Cov}(f_2, e_i)) \approx 0$ and the Fama-MacBeth slope is close to $E[r_p]$.

When the test assets are sorted on a characteristic related to the omitted factor f_2, the residuals e_i inherit systematic exposure to f_2 that varies across the sort, and if this exposure correlates with $\beta_{i,p}$, the OFB can be large.

When the test assets are sorted on a characteristic unrelated to both the factor and the omitted component, the betas $\beta_{i,p}$ may cluster tightly (low $\mathrm{Var}_{\mathrm{cs}}(\beta_{i,p})$), amplifying the ratio in (14) and producing noisy or extreme estimates of γ_1.

2.7 The Sharpe ratio constrains the budget, not the distribution

From the variance decomposition of the SDF projection (1):

$$\mathrm{Var}(m) = b^2\ \mathrm{Var}(r_p) + \mathrm{Var}(\varepsilon).$$

Substituting $b = -E[m]\, E[r_p]/\,\mathrm{Var}(r_p)$ and dividing by $E[m]^2$:

$$\frac{\mathrm{Var}(m)}{E[m]^2} = \frac{E[r_p]^2}{\mathrm{Var}(r_p)} + \frac{\mathrm{Var}(\varepsilon)}{E[m]^2} = \mathrm{SR}^2(p) + \frac{\mathrm{Var}(\varepsilon)}{E[m]^2}. \tag{16}$$

When m is the minimum-variance SDF (the projection of any valid SDF onto the asset return space), the left-hand side equals $\mathrm{SR}^2(\text{tangency})$ by the Hansen-Jagannathan (1991) bound. Therefore:

$$\mathrm{SR}^2(\text{tangency}) = \mathrm{SR}^2(p) + \frac{\mathrm{Var}(\varepsilon)}{E[m]^2}. \tag{17}$$

A higher $\mathrm{SR}(p)$ reduces $\mathrm{Var}(\varepsilon)/E[m]^2$, the aggregate pricing error from using p alone. $\mathrm{SR}(p)$ measures how much of the SDF's pricing power p captures.

However, the OFB in (5) depends on $\mathrm{Cov}_{\mathrm{cs}}(\beta_{i,p}, \mathrm{Cov}(\varepsilon, r_i))$, not on $\mathrm{Var}(\varepsilon)$. The Sharpe ratio constrains the total size of the ε residual but not how the individual $\mathrm{Cov}(\varepsilon, r_i)$ terms are distributed across assets with different betas. The same total $\mathrm{Var}(\varepsilon)$ can produce zero OFB (if mispricing is orthogonal to betas) or large OFB (if mispricing concentrates in high- or low-beta assets).

Remark 2. *The Sharpe ratio of a portfolio constrains the aggregate mispricing from using that portfolio as a single factor, but does not by itself determine the sign or magnitude of the cross-sectional price of risk estimated in a Fama-MacBeth regression. For any* $\mathrm{SR}(p) < \mathrm{SR}(\textit{tangency})$*, the OFB can take any value, including values that flip the sign of* γ_1.

In the limit $\mathrm{SR}(p) \to \mathrm{SR}(\text{tangency})$, $\mathrm{Var}(\varepsilon) \to 0$. Since $|\,\mathrm{Cov}(\varepsilon, r_i)|^2 \leq \mathrm{Var}(\varepsilon)\ \mathrm{Var}(r_i)$ by the Cauchy-Schwarz inequality, each individual $|\,\mathrm{Cov}(\varepsilon, r_i)| \to 0$, and therefore the OFB in (5) vanishes: $\gamma_1 \to E[r_p]$. But short of spanning the SDF, the cross-sectional distribution of the residual covariances is unconstrained by the Sharpe ratio.

2.8 Inclusion in the SDF versus projection

The preceding results clarify a distinction the literature often conflates.

Suppose first that r_p is a component of the SDF in the sense that m is an affine function of r_p (possibly among other factors). In the single-factor case, this means

$$m = a + b\, r_p, \tag{18}$$

with no residual: $\varepsilon = 0$. Then the pricing equation (4) reduces to

$$E[r_i] = \beta_{i,p}\, E[r_p] \qquad \text{for all } i,$$

and the Fama-MacBeth slope is $\gamma_1 = E[r_p]$ exactly. The Sharpe ratio of p equals the tangency Sharpe ratio, and the cross-sectional slope coincides with the factor's expected return. The equivalence between discount factor and beta pricing representations requires exactly this condition: if the SDF is linear in r_p and prices all assets, then expected returns are linear in betas and the slope is $E[r_p]$.

In general, the researcher does not know whether r_p is in the SDF or merely correlated with it. The projection (1) is always valid, for any portfolio p, without assuming that p is a component of m. The residual ε captures everything in the SDF that r_p does not span. When $\varepsilon \neq 0$, the pricing equation (4) contains the additional term $\mathrm{Cov}(\varepsilon, r_i)/E[m]$, and the Fama-MacBeth slope deviates from $E[r_p]$ by the OFB in (5).

Remark 3. *SDF spanning is sufficient but not necessary for the Fama-MacBeth slope to coincide with the factor's expected return.*

(a) *If $\varepsilon = 0$ (equivalently, $\mathrm{SR}(p) = \mathrm{SR}(\textit{tangency})$), then $\gamma_1 = E[r_p]$ regardless of the test assets.*

(b) *If $\varepsilon \neq 0$ (the factor does not span the SDF), $\gamma_1 = E[r_p]$ still holds provided*

$$\mathrm{Cov}_{\mathrm{cs}}(\beta_{i,p},\ \mathrm{Cov}(\varepsilon, r_i)) = 0,$$

that is, the pricing contributions of the omitted SDF components are cross-sectionally uncorrelated with the factor betas.

Part (a) is the standard equivalence between discount factor and beta pricing representations. For γ_1 to equal $E[r_p]$ requires only the weaker condition in part (b): the omitted SDF components need not be absent, only balanced across beta groups.

Three conditions of increasing strength govern the relationship between a factor and the SDF. First, the univariate projection coefficient b in (2) is nonzero: since $b = -E[m]\,E[r_p]/\operatorname{Var}(r_p)$, this holds whenever $E[r_p] \neq 0$, a condition satisfied by every long-short factor with a nonzero mean. The bar for "being in the SDF" in the univariate sense is trivially low. Second, the factor has marginal pricing power in the true SDF: $b_p \neq 0$ in the multivariate projection $m = a + b'f^* + \varepsilon^*$, where f^* collects all priced factors. This means the factor contributes pricing information beyond all other factors. Third, the factor set spans the SDF: $\varepsilon = 0$, so the factors capture the entire pricing kernel. The first condition carries no economic content beyond a nonzero expected return. The third is empirically verifiable given a model (via the GRS test). The second is untestable without knowing the complete set of priced factors, the cross-sectional analog of Roll (1977): whether a particular factor belongs to the true SDF requires observing the true SDF.

A factor can satisfy the first condition while being entirely redundant given the true factors. Suppose the true SDF is $m = a_0 + b_1 f_1 + b_2 f_2$ and the candidate factor is $r_p = \frac{1}{2}f_1 + \frac{1}{2}f_2 + \nu$, where ν is idiosyncratic noise. The factor is spanned by f_1 and f_2: its time-series alpha with respect to the true model is zero, and $E[r_p] = \frac{1}{2}E[f_1] + \frac{1}{2}E[f_2]$ reflects only its loadings on the true factors. Yet the univariate projection $m = a + b\,r_p + \varepsilon$ yields $b \neq 0$ (because $E[r_p] \neq 0$) and $\varepsilon \neq 0$ (because r_p combines the true factors with the wrong weights and adds noise). The single-factor Fama-MacBeth regression produces $\gamma_1 \neq E[r_p]$ in general. Being "spanned by" the true factors and "spanning" the SDF are opposite conditions on opposite objects, and the first does not preclude the second from failing.

The gap between parts (a) and (b) is the paper's central distinction. SDF spanning means the factor is included in the discount factor; the projection (1) is valid for any portfolio whether or not it is included. The Fama-MacBeth regression operates through the projection, so γ_1 can equal $E[r_p]$ even when the factor does not span the SDF, as long as the omitted SDF components do not load differentially on high- versus low-beta assets. Economically, the orthogonality condition in (b) holds when the unpriced SDF variation affects all beta-sorted groups symmetrically: for instance, if the omitted factor represents an industry-specific risk that is balanced across high- and low-beta portfolios. Conversely, the condition fails when the omitted factor loads primarily on assets at one end of the beta distribution, as when a missing momentum or liquidity factor covaries more strongly with high-beta growth stocks than with low-beta value stocks.

The Sharpe ratio and the Fama-MacBeth slope answer different questions. The Sharpe ratio asks: how much of the SDF does this portfolio capture? The Fama-MacBeth slope asks: in the cross-section of test assets, does variation in betas on this portfolio explain variation in expected returns? The first is a property of the portfolio alone. The second depends on the interaction between the portfolio, the unspanned SDF component, and the test assets. A high Sharpe ratio does not guarantee that the cross-sectional price of beta equals the factor's expected return.

2.9 Pricing errors and the Sharpe ratio budget

The decomposition (5) uses the SDF residual ε. Restating γ_1 in terms of population pricing errors (alphas) connects to the time-series regression framework and the Gibbons, Ross, and Shanken (1989) bound.

From the time-series regression (12), the population pricing error of asset i with respect to factor p is $\alpha_{i,p} = E[r_i] - \beta_{i,p}\, E[r_p]$. Substituting $E[r_i] = \beta_{i,p}\, E[r_p] + \alpha_{i,p}$ into the cross-sectional OLS formula for γ_1:

$$\gamma_1 = E[r_p] + \frac{\mathrm{Cov}_{\mathrm{cs}}(\beta_{i,p},\ \alpha_{i,p})}{\mathrm{Var}_{\mathrm{cs}}(\beta_{i,p})}. \tag{19}$$

Equation (19) restates the two components of γ_1: the factor's expected return $E[r_p]$ and the cross-sectional regression coefficient of the alphas on the betas. The second component has the structure of an omitted variable bias, but the omitted "variable" is the unspanned SDF component, which no finite factor model includes. This component is not a finite-sample artifact or an estimation choice: it is a structural feature of any cross-sectional regression that does not span the complete pricing kernel. The expression is algebraically equivalent to (5), since the pricing equation (4) implies $\alpha_{i,p} = -\operatorname{Cov}(\varepsilon, r_i)/E[m]$, but it makes the economic content more transparent. If high-beta assets tend to have positive alphas (the factor underprices them), γ_1 exceeds $E[r_p]$. If high-beta assets tend to have negative alphas, as in the betting-against-beta phenomenon, γ_1 falls below $E[r_p]$. The two components offset exactly when the alphas are cross-sectionally uncorrelated with betas.

The alphas and the SDF residual have a concrete interpretation in terms of the tangency portfolio. A risk-free rate R_f exists, so $E[m] = 1/R_f$. Let r^* denote the tangency (maximum Sharpe ratio) portfolio of the asset universe. The minimum-variance SDF is

$$m^* = \frac{1}{R_f}\left(1 - \frac{E[r^*]}{\operatorname{Var}(r^*)}\,(r^* - E[r^*])\right), \tag{20}$$

which is affine in r^*: $m^* = a^* + b^*\, r^*$ with $b^* = -E[r^*]/(R_f \operatorname{Var}(r^*))$. For any portfolio p, project the tangency portfolio onto r_p:

$$r^* = \alpha_p^* + \beta_p^*\, r_p + \eta_p, \qquad \operatorname{Cov}(\eta_p, r_p) = 0, \tag{21}$$

where η_p is the component of the tangency portfolio that p does not span. Substituting into $m^* = a^* + b^*\, r^*$ yields

$$m^* = \underbrace{(a^* + b^*\alpha_p^*)}_{a} \;+\; \underbrace{b^*\beta_p^*}_{b}\, r_p \;+\; \underbrace{b^*\eta_p}_{\varepsilon},$$

which is the projection of m^* onto r_p (since $\operatorname{Cov}(\eta_p, r_p) = 0$ implies $\operatorname{Cov}(\varepsilon, r_p) = 0$). Comparing with (1):

$$\varepsilon = b^*\,\eta_p. \tag{22}$$

The SDF residual from using p as a single factor is proportional to η_p, the component of the tangency portfolio that p fails to capture. When $p = r^*$, $\eta_p = 0$ and $\varepsilon = 0$; otherwise, the residual inherits the properties of the missing tangency component.

The pricing errors take a correspondingly explicit form. From (4) and (22), $\alpha_{i,p} = -\operatorname{Cov}(\varepsilon, r_i)/E[m] = -R_f\, b^* \operatorname{Cov}(\eta_p, r_i)$, so

$$\alpha_{i,p} = \frac{E[r^*]}{\operatorname{Var}(r^*)} \operatorname{Cov}(\eta_p, r_i). \tag{23}$$

Each asset's pricing error is proportional to its covariance with the missing tangency component. The Sharpe ratio of p constrains $\operatorname{Var}(\eta_p)$ through (17), but not the cross-sectional distribution of $\operatorname{Cov}(\eta_p, r_i)$ across assets with different betas.

The Sharpe ratio constrains the alphas through the quadratic form in (24), not the inner product that governs the FM bias. For a fixed set of N test assets, the Gibbons, Ross, and Shanken (1989) decomposition states:

$$\mathrm{SR}^2(\text{tangency}) = \mathrm{SR}^2(p) + \alpha' \Sigma_e^{-1} \alpha, \tag{24}$$

where $\alpha = (\alpha_{1,p}, \ldots, \alpha_{N,p})'$ is the vector of pricing errors and Σ_e is the covariance matrix of the time-series residuals e_i from (12). A higher $\mathrm{SR}(p)$ reduces the right-hand side, shrinking the quadratic form $\alpha' \Sigma_e^{-1} \alpha$ toward zero. In the limit $\mathrm{SR}(p) = \mathrm{SR}(\text{tangency})$, all alphas are zero, the OFB component vanishes, and $\gamma_1 = E[r_p]$. The question is whether, short of spanning, the Sharpe ratio constrains γ_1 tightly enough to pin down its sign and magnitude.

Proposition 1 (Cauchy-Schwarz bound)**.** *Applying the Cauchy-Schwarz inequality with the Σ_e^{-1} inner product to $\tilde{\beta}'\alpha$, where $\tilde{\beta}_i = \beta_{i,p} - \bar{\beta}_p$, and substituting* (24)*:*

$$|\gamma_1 - E[r_p]| \;\leq\; \frac{\sqrt{\tilde{\beta}' \Sigma_e\, \tilde{\beta}}}{\tilde{\beta}'\tilde{\beta}} \sqrt{\mathrm{SR}^2(\textit{tangency}) - \mathrm{SR}^2(p)}. \tag{25}$$

The bound tightens as $\mathrm{SR}(p) \to \mathrm{SR}(\textit{tangency})$, but its rate of tightening depends on the prefactor $\tilde{\beta}'\Sigma_e\,\tilde{\beta}/(\tilde{\beta}'\tilde{\beta})^2$: the residual variance concentrated in assets with extreme betas. In typical cross-sections, sorted portfolios at the extremes of a characteristic have the highest residual risk, so the assets that drive the FM regression are also the assets where the factor *model fits worst. Section 4.6 confirms that the bound permits wrong-sign slopes for 83–100% of factors across all plausible tangency Sharpe ratios, making it loose in practice.*

Remark 4 (GLS bound)**.** *The GLS slope* (11) *admits a tighter bound than its OLS counterpart. Applying Cauchy-Schwarz directly in the* Σ_e^{-1} *inner product:*

$$|\gamma_1^{\text{GLS}} - E[r_p]| \;\le\; \frac{\sqrt{\text{SR}^2(\textit{tangency}) - \text{SR}^2(p)}}{\sqrt{\tilde{\beta}'\Sigma_e^{-1}\tilde{\beta}}}. \tag{26}$$

By a further application of Cauchy-Schwarz, $(\tilde{\beta}'\tilde{\beta})^2 \le (\tilde{\beta}'\Sigma_e\,\tilde{\beta})(\tilde{\beta}'\Sigma_e^{-1}\tilde{\beta})$*, which implies that the GLS prefactor is weakly smaller than the OLS prefactor from Proposition 1, with equality when* Σ_e *is proportional to the identity. The bound is tighter because the GLS estimator and the GRS quadratic form share the same* Σ_e^{-1} *metric, eliminating the metric mismatch that loosens the OLS bound. The improvement is quantitative, not qualitative: the bound* (26) *still permits* $|\gamma_1^{\text{GLS}} - E[r_p]| \ge |E[r_p]|$*, and hence a zero or wrong-sign GLS slope, whenever* $\text{SR}^2(\textit{tangency}) - \text{SR}^2(p) \ge E[r_p]^2 \cdot \tilde{\beta}'\Sigma_e^{-1}\tilde{\beta}$*. The Sharpe ratio disciplines GLS more tightly than OLS, but not enough to guarantee the correct sign or magnitude of the estimated price of risk.*

Remark 5 (Data-based bound via the GRS statistic)**.** *Substituting* $\text{SR}^2(\textit{tangency}) - \text{SR}^2(p) = \alpha'\Sigma_e^{-1}\alpha$ *into* (25) *eliminates the unobservable tangency Sharpe ratio:*

$$|\gamma_1 - E[r_p]| \;\le\; \frac{\sqrt{\tilde{\beta}'\Sigma_e\,\tilde{\beta}}}{\tilde{\beta}'\tilde{\beta}}\,\sqrt{\alpha'\Sigma_e^{-1}\alpha}. \tag{27}$$

The right-hand side is estimable from data: $\alpha'\Sigma_e^{-1}\alpha$ *is the numerator of the Gibbons, Ross, and Shanken (1989) test statistic (up to a degrees-of-freedom scalar), and the prefactor involves only betas and residual variances from the time-series regressions. The bound does not require knowledge of the tangency Sharpe ratio. It is nevertheless uninformative for the same reason as* (25)*:* $\alpha'\Sigma_e^{-1}\alpha$ *constrains the norm of the alpha vector, not its direction relative to beta. The same GRS statistic is consistent with* $\gamma_1 = E[r_p]$ *(alphas orthogonal to betas)* *and with $|\gamma_1 - E[r_p]|$ at the Cauchy-Schwarz maximum (alphas perfectly aligned with betas). The direct sample estimate $\widehat{\mathrm{Cov}}_{\mathrm{cs}}(\hat{\beta}, \hat{\alpha})/\widehat{\mathrm{Var}}_{\mathrm{cs}}(\hat{\beta})$ is computable from the same time-series regressions and is strictly more informative, because it preserves the directional information that the norm-based bound discards.*

Example: the CAPM and the flat security market line. Setting $p = r_m$ (the market portfolio) recovers the CAPM as a special case. The tangency decomposition (21) becomes $r^* = \alpha_m^* + \beta_m^* r_m + \eta_m$, where η_m is the component of the tangency portfolio orthogonal to the market, and the CAPM alphas satisfy $\alpha_{i,m} = \frac{E[r^*]}{\mathrm{Var}(r^*)} \mathrm{Cov}(\eta_m, r_i)$ by (23). The population FM slope from a cross-sectional regression of expected returns on market betas is

$$\gamma_1 = E[r_m] + \frac{\mathrm{Cov}_{\mathrm{cs}}(\beta_{i,m},\ \alpha_{i,m})}{\mathrm{Var}_{\mathrm{cs}}(\beta_{i,m})}.$$

The well-documented finding that the security market line is flatter than the CAPM predicts (Black 1972; Fama and French 1992) corresponds to $\gamma_1 < E[r_m]$, requiring $\mathrm{Cov}_{\mathrm{cs}}(\beta_{i,m}, \alpha_{i,m}) < 0$: high-beta assets tend to have negative CAPM alphas. The resulting flat security market line is the betting-against-beta pattern (Frazzini and Pedersen 2014). The decomposition provides a structural explanation: η_m, the non-market component of the tangency portfolio (reflecting size, value, momentum, and other priced factors), covaries with test assets in a pattern that is negatively correlated with market betas. The cross-sectional price of market beta, γ_1, is near zero not because the market earns a low return but because the OFB approximately offsets $E[r_m]$. The market's Sharpe ratio (≈ 0.4 annualized; Cochrane 2005) is well below the tangency Sharpe ratio, so the alpha budget (24) imposes little constraint on the cross-sectional covariance that determines the second component.

2.10 The Fama-MacBeth intercept and the zero-beta rate

The OFB affects the slope and the intercept of the cross-sectional regression simultaneously. The population Fama-MacBeth regression is

$$E[r_i] = \gamma_0 + \gamma_1\, \beta_{i,p},$$

where r_i denotes excess returns. By the OLS intercept formula, $\gamma_0 = \bar{\mu} - \gamma_1\, \bar{\beta}$, where $\bar{\mu} = N^{-1} \sum_i E[r_i]$ and $\bar{\beta} = N^{-1} \sum_i \beta_{i,p}$ are the cross-sectional means. Substituting $E[r_i] = \beta_{i,p}\, E[r_p] + \alpha_{i,p}$ from (4) gives $\bar{\mu} = \bar{\beta}\, E[r_p] + \bar{\alpha}$, where $\bar{\alpha} = N^{-1} \sum_i \alpha_{i,p}$ is the average pricing error. Combining with $\gamma_1 = E[r_p] + \text{OFB}$ from (19):

$$\gamma_0 = \bar{\alpha} - \text{OFB} \cdot \bar{\beta}. \tag{28}$$

Under correct specification ($\alpha_{i,p} = 0$ for all i), $\bar{\alpha} = 0$ and $\text{OFB} = 0$, so $\gamma_0 = 0$: the intercept of an excess-return cross-sectional regression is zero, as it should be. Under misspecification, the intercept absorbs the OFB through the cross-sectional centroid constraint. In total returns, the intercept becomes $\gamma_0 = R_f + \bar{\alpha} - \text{OFB} \cdot \bar{\beta}$, and $\gamma_0 = R_f$ under correct specification.

The intercept and slope are linked through a seesaw. The regression line passes through the centroid $(\bar{\beta}, \bar{\mu})$, so any OFB that pushes γ_1 below $E[r_p]$ must push γ_0 above its null value (when $\bar{\beta} > 0$), and vice versa. Equation (28) makes the transfer explicit: each unit of negative OFB lowers the slope by one unit and raises the intercept by $\bar{\beta}$ units. The two deviations are not independent anomalies; they are one anomaly viewed at two points on the regression line.

The zero-beta rate puzzle as omitted factor bias. In the CAPM context ($p = r_m$), Black (1972) and Fama and MacBeth (1973) documented that the Fama-MacBeth intercept exceeds the risk-free rate ($\gamma_0 > R_f$ in total returns) while the slope falls below the equity premium ($\gamma_1 < E[r_m]$). Black (1972) proposed the zero-beta CAPM as an explanation: with no risk-free asset, the intercept estimates the expected return on a zero-beta portfolio, which may exceed R_f. The OFB formula provides an alternative explanation that does not require abandoning the risk-free asset. In total returns with $\bar{\beta} \approx 1$ (the average market beta of test assets):

$$\gamma_0 - R_f \;\approx\; \bar{\alpha} + |\text{OFB}|,$$

since the market's OFB is negative ($\text{Cov}_{\text{cs}}(\beta_m, \alpha_m) < 0$). The "too high" intercept reflects the negative OFB transferring expected return from the slope to the intercept. The zero-beta rate $1/E[m]$ may equal R_f exactly; the Fama-MacBeth intercept overshoots R_f because the cross-sectional regression misattributes part of the expected return, allocating too little to the slope (due to the betting-against-beta pattern in the alphas) and too much to the intercept. The "too high zero-beta rate" and the "too flat security market line" are one finding, not two.

The OFB explanation and Black's zero-beta CAPM are observationally equivalent in a single cross-section with a single factor: both produce $\gamma_0 > R_f$ and $\gamma_1 < E[r_m]$. They differ in what generates the elevated intercept. Under the zero-beta CAPM, γ_0 estimates a genuine economic quantity (the expected return on a zero-beta portfolio, elevated by borrowing constraints), and a correctly specified model should produce a stable γ_0 regardless of the test assets or the number of factors. Under the OFB explanation, γ_0 is a regression artifact: it absorbs the OFB through the centroid constraint, so it should vary with test-asset choice (as the OFB does) and converge toward R_f as additional factors reduce the unspanned SDF component. The two explanations are therefore distinguishable in principle: if adding factors to the model moves γ_0 toward R_f and the intercept varies across test-asset configurations, the OFB explanation is operative. The OFB explanation does not rule out a genuine zero-beta premium; it shows that the Fama-MacBeth intercept conflates any true zero-beta premium with the mechanical seesaw from the OFB, and the latter can be quantified directly from equation (28).

2.11 When does a high Sharpe ratio determine γ_1?

A high Sharpe ratio shrinks the aggregate mispricing budget but does not by itself determine the Fama-MacBeth slope. Under what conditions does a high Sharpe ratio pin down γ_1? This subsection answers by decomposing the second component of γ_1 into a magnitude that the Sharpe ratio controls and a directional term that it does not.

Decomposition into magnitude and direction. The OFB from (19) is the cross-sectional regression slope of $\alpha_{i,p}$ on $\beta_{i,p}$. Rewriting it in terms of the cross-sectional correlation and standard deviations:

$$\gamma_1 - E[r_p] = \rho_{\text{cs}}(\beta_{i,p},\, \alpha_{i,p}) \cdot \frac{\text{SD}_{\text{cs}}(\alpha_{i,p})}{\text{SD}_{\text{cs}}(\beta_{i,p})}, \tag{29}$$

where ρ_{cs} and SD_{cs} denote the cross-sectional correlation and standard deviation over the N test assets.

The two components play distinct roles. The ratio $\text{SD}_{\text{cs}}(\alpha)/\text{SD}_{\text{cs}}(\beta)$ measures the magnitude of mispricing relative to the dispersion in betas. The Sharpe ratio constrains this term through the GRS bound (24). Under homogeneous residual variances ($\Sigma_e = \sigma_e^2\, I_N$), the bound becomes

$$\text{SD}_{\text{cs}}(\alpha) \;\leq\; \frac{\sigma_e}{\sqrt{N}} \sqrt{\text{SR}^2(\text{tangency}) - \text{SR}^2(p)}, \tag{30}$$

so a high $\text{SR}(p)$ reduces the maximal alpha dispersion. The bound uses $\text{SD}_{\text{cs}}(\alpha)$ rather than $\|\alpha\|/\sqrt{N}$ because the cross-sectional regression includes an intercept that absorbs $\bar{\alpha}$.[3] The cross-sectional correlation $\rho_{\text{cs}}(\beta, \alpha)$, by contrast, measures the directional alignment between betas and pricing errors: do high-beta assets tend to be the mispriced ones? This term ranges from -1 to $+1$ and is not constrained by the Sharpe ratio of p.

3. The general version with heterogeneous Σ_e replaces σ_e^2/N with the appropriate eigenvalue of Σ_e, but the qualitative conclusion is unchanged.

Corollary 1. *The alignment term decomposes as*

$$\gamma_1 - E[r_p] = \rho_{\text{cs}}(\beta_{i,p}, \alpha_{i,p}) \cdot \frac{\text{SD}_{\text{cs}}(\alpha_{i,p})}{\text{SD}_{\text{cs}}(\beta_{i,p})},$$

where $\text{SD}_{\text{cs}}(\alpha)/\text{SD}_{\text{cs}}(\beta)$ *is the magnitude of mispricing relative to the dispersion in betas, and* $\rho_{\text{cs}}(\beta, \alpha)$ *is the directional alignment between betas and pricing errors. The Sharpe ratio constrains the magnitude through the GRS bound* (24)*: a higher* $\text{SR}(p)$ *reduces the maximal* $\text{SD}_{\text{cs}}(\alpha)$*. But the Sharpe ratio does not constrain* $\rho_{\text{cs}}(\beta, \alpha)$*, which ranges freely from* -1 *to* $+1$ *for any* $\text{SR}(p) < \text{SR}(\textit{tangency})$*. The alignment term vanishes if either:*

(a) $\text{SD}_{\text{cs}}(\alpha) = 0$ *(correct specification, governed by the Sharpe ratio); or*

(b) $\rho_{\text{cs}}(\beta_{i,p}, \alpha_{i,p}) = 0$ *(orthogonal mispricing, not governed by the Sharpe ratio).*

Condition (b) can hold even when the model is misspecified, if the mispricing happens not to load differentially on high- versus low-beta assets. Conversely, condition (b) can fail even when the Sharpe ratio is high, if the residual mispricing is concentrated in assets with extreme betas. A direct consequence is the condition for a wrong-sign slope.

Corollary 2. *The Fama-MacBeth slope has the wrong sign (*$\gamma_1 \leq 0$ *when* $E[r_p] > 0$*) if and only if*

$$\rho_{\text{cs}}(\beta_{i,p}, \alpha_{i,p}) \leq -\frac{E[r_p] \cdot \text{SD}_{\text{cs}}(\beta_{i,p})}{\text{SD}_{\text{cs}}(\alpha_{i,p})}.$$

The required directional misalignment is smaller when $E[r_p]$ *is small relative to alpha dispersion, the regime where most factors reside.*

Remark 6 (Scale invariance)**.** *Both the Sharpe ratio and the wrong-sign condition are invariant to the scale of the factor portfolio. Replacing* r_p *with* $k\,r_p$ *for* $k > 0$ *sends* $E[r_p] \to k\,E[r_p]$*,* $\sigma(r_p) \to k\,\sigma(r_p)$*,* $\beta_{i,p} \to \beta_{i,p}/k$*,* $\alpha_{i,p} \to \alpha_{i,p}$*,* $\gamma_1 \to k\,\gamma_1$*, and* $\text{OFB} \to k\,\text{OFB}$*, so* $\text{SR}(p)$*,* $\text{OFB}/E[r_p]$*, and* $\text{sign}(\gamma_1)$ *are unchanged. In particular, the wrong-sign condition can be expressed without reference to the factor's expected return or its Sharpe ratio magnitude.*

Define $b_i = \text{Cov}(r_i, r_p)/\sigma(r_p)$ *(scale-invariant) and* $D^* = \text{Cov}_{\text{cs}}(b_i, E[r_i])/\operatorname{Var}_{\text{cs}}(b_i)$. *Then* $\gamma_1/E[r_p] = D^*/\text{SR}(p)$, *and*

$$\gamma_1 \cdot E[r_p] < 0 \quad \Longleftrightarrow \quad \text{sign}(D^*) \neq \text{sign}(E[r_p]).$$

The quantity D^* *measures whether assets whose returns covary more with the factor tend to have higher expected returns; it is a structural property of the cross-section that does not depend on the factor's scale. A declining wrong-sign rate across Sharpe-ratio quintiles therefore reflects that high-Sharpe-ratio factors are better aligned with the cross-section of expected returns, not a mechanical threshold from larger* $|E[r_p]|$.

The decomposition provides a geometric interpretation that complements the ellipsoid picture from Proposition 1. The Sharpe ratio constrains the alpha vector to a shrinking ellipsoid; a high Sharpe ratio makes the vector short. But the OFB depends on the projection of this vector onto the beta vector, which is the product of the vector's length and the cosine of the angle between them. Even a short alpha vector, if it points along the beta vector ($|\rho| \approx 1$), produces a large projection per unit length. A high Sharpe ratio shrinks the norm of α; it cannot rotate α away from β.

The tangency residual and cross-sectional beta alignment. The tangency portfolio representation makes the directional component concrete. From (23), the pricing error of asset i is

$$\alpha_{i,p} = \frac{E[r^*]}{\text{Var}(r^*)} \, \text{Cov}(\eta_p, r_i),$$

where η_p is the component of the tangency portfolio orthogonal to r_p (equation (21)). Define the *residual beta*

$$\beta_{i,\eta} = \frac{\text{Cov}(\eta_p, r_i)}{\text{Var}(\eta_p)}, \tag{31}$$

the beta of asset i on the tangency residual. Then $\alpha_{i,p} = c_p \, \beta_{i,\eta}$, where $c_p = E[r^*] \, \operatorname{Var}(\eta_p)/\operatorname{Var}(r^*)$ is a positive scalar that depends on the tangency portfolio and vanishes if and only if

$\mathrm{Var}(\eta_p) = 0$, which is equivalent to $\mathrm{SR}(p) = \mathrm{SR}(\text{tangency})$.

Substituting into the OFB formula (19):

$$\gamma_1 - E[r_p] = c_p \cdot \frac{\mathrm{Cov}_{\mathrm{cs}}(\beta_{i,p},\, \beta_{i,\eta})}{\mathrm{Var}_{\mathrm{cs}}(\beta_{i,p})} = c_p \cdot \delta_{\eta|p}, \tag{32}$$

where $\delta_{\eta|p}$ is the cross-sectional regression coefficient of the residual betas $\beta_{i,\eta}$ on the factor betas $\beta_{i,p}$. The OFB is the product of two terms: c_p, which measures the factor's unspanned pricing power (governed by the Sharpe ratio gap), and $\delta_{\eta|p}$, which measures whether the test assets' exposures to the missing SDF component align with their exposures to the included factor (a property of the cross-section that the Sharpe ratio does not constrain).

The constant c_p admits a direct expression in terms of the Sharpe ratio gap. Since $\mathrm{Var}(\eta_p) = \mathrm{Var}(r^*)\left[\mathrm{SR}^2(\text{tangency}) - \mathrm{SR}^2(p)\right]/\mathrm{SR}^2(\text{tangency})$,

$$c_p = E[r^*] \cdot \frac{\mathrm{SR}^2(\text{tangency}) - \mathrm{SR}^2(p)}{\mathrm{SR}^2(\text{tangency})}. \tag{33}$$

As $\mathrm{SR}(p) \to \mathrm{SR}(\text{tangency})$, $c_p \to 0$ and the OFB vanishes regardless of $\delta_{\eta|p}$. But for any fixed gap $\mathrm{SR}(p) < \mathrm{SR}(\text{tangency})$, the magnitude of the OFB is determined by $\delta_{\eta|p}$, which can take any value.

Time-series orthogonality does not imply cross-sectional orthogonality. The factors r_p and η_p are uncorrelated as random variables: $\mathrm{Cov}(r_p, \eta_p) = 0$ by construction of the tangency projection (21). Separate univariate regressions therefore correctly estimate $\beta_{i,p}$ and $\beta_{i,\eta}$: there is no time-series omitted variable problem. But time-series orthogonality does not imply that $\mathrm{Cov}_{\mathrm{cs}}(\beta_{i,p}, \beta_{i,\eta}) = 0$. The cross-sectional covariance between betas on orthogonal factors is determined by how assets distribute their loadings across factors, not by the factors' time-series properties.

The distinction is between a property of the factors (time-series uncorrelatedness) and a property of the cross-section of assets (how loadings co-vary). Factors do not choose which assets load on them; assets load on factors according to their economic characteristics, and those characteristics typically generate exposure to multiple systematic risks simultaneously. A small-cap value stock may have high market beta and high exposure to the tangency residual, not because these factors are time-series correlated but because the firm's characteristics place it in a region of the cross-section where multiple risk exposures are elevated.

Why the directional component is generically nonzero. The cross-sectional regression coefficient $\delta_{\eta|p}$ is zero only when the test assets' exposures to the factor and to the tangency residual are cross-sectionally uncorrelated. This condition requires a specific cancellation: the assets with high $\beta_{i,p}$ must, on average, have the same $\beta_{i,\eta}$ as the assets with low $\beta_{i,p}$.

When there is a single omitted factor with expected return μ_2 and betas $\beta_{i,2}$, the alpha is $\alpha_{i,p} = \beta_{i,2}\,\mu_2$ (with independent factors), and $\rho_{\text{cs}}(\beta_{i,p}, \alpha_{i,p}) = \rho_{\text{cs}}(\beta_{i,p}, \beta_{i,2}) \cdot \text{sign}(\mu_2)$. The cross-sectional correlation between the included and omitted betas is a fixed property of the economy, not a random draw that might happen to be zero. When there are $K-1$ omitted factors, the condition becomes

$$\sum_{k=2}^{K} \mu_k \ \text{Cov}_{\text{cs}}(\beta_{i,p},\ \beta_{i,k}) = 0, \tag{34}$$

which requires exact cancellation among the omitted factors' contributions. Each term $\mu_k\ \text{Cov}_{\text{cs}}(\beta_{i,p}, \beta_{i,k})$ reflects how the k-th omitted factor's risk premium, weighted by the cross-sectional alignment of its betas with the included factor's betas, contributes to the OFB. For this sum to vanish, positive and negative contributions must offset precisely, a knife-edge condition that holds only for special configurations of risk premia and beta structures. Rewriting (34) as a correlation gives ρ in terms of economic primitives:

$$\rho_{\text{cs}}(\beta_{i,p}, \alpha_{i,p}) = \frac{\sum_{k\neq p} \mu_k\ \text{Cov}_{\text{cs}}(\beta_{i,p}, \beta_{i,k})}{\text{SD}_{\text{cs}}(\beta_{i,p}) \cdot \text{SD}_{\text{cs}}(\alpha_{i,p})}. \tag{35}$$

The sign of ρ is the sign of the risk-premium-weighted sum of cross-sectional beta covariances. The directional component is negative whenever the omitted factors whose betas are negatively correlated with $\beta_{i,p}$ across assets carry, in aggregate, larger risk premia than those whose betas are positively correlated.

Economically, the cancellation fails because the same characteristics that drive $\beta_{i,p}$ also influence the omitted betas $\beta_{i,k}$. If r_p is the market portfolio, the tangency residual η_p captures size, value, profitability, and other priced dimensions. A firm's market beta reflects its cash flow duration, leverage, and operating risk, characteristics that also determine its exposure to these omitted factors. The CAPM example illustrates this pattern: the negative $\mathrm{Cov}_{\mathrm{cs}}(\beta_{i,m}, \alpha_{i,m})$ that produces the flat security market line reflects the empirical pattern that high-market-beta stocks (typically growth firms with long-duration cash flows) have low exposure to the value and profitability components of the tangency residual. The correlation $\rho_{\mathrm{cs}}(\beta_m, \alpha_m)$ is not an accident; it is a structural feature of the cross-section.

The same logic applies to any factor. A momentum factor's betas correlate with its omitted value exposure across test assets because both reflect duration and reversal characteristics. A profitability factor's betas correlate with its omitted investment exposure because the same accounting fundamentals drive both. The cross-sectional correlation $\rho_{\mathrm{cs}}(\beta_{i,p}, \alpha_{i,p})$ is generically nonzero not because of bad luck in test-asset selection but because the economic determinants of factor exposures do not respect the orthogonal decomposition of the SDF.

Why ρ is predominantly negative. The tangency residual representation makes the sign prediction concrete. From (32), $\alpha_{i,p} = c_p \, \beta_{i,\eta}$ with $c_p > 0$, so $\rho_{\mathrm{cs}}(\beta_{i,p}, \alpha_{i,p}) = \rho_{\mathrm{cs}}(\beta_{i,p}, \beta_{i,\eta})$. The tangency portfolio is the optimal diversified combination of all priced factors; its residual η_p captures the balanced, multi-dimensional pricing that r_p alone misses. Assets with extreme $\beta_{i,p}$ are specialists concentrated in one corner of the characteristic space. The tangency residual rewards diversified exposure across the orthogonal priced dimensions that these specialists lack: a deep value stock has high exposure to the value factor but typically low exposure to momentum and profitability, the dimensions η_p captures. High-β_p assets therefore tend to have low $\beta_{i,\eta}$, producing $\rho < 0$.

The same pattern follows from (35). Most long-short factors carry positive risk premia, and characteristic-sorted factor betas exhibit a substitution structure: stocks sorted to the extreme on one characteristic tend to load less on other priced factors whose sorts emphasize different characteristics. When most $\mu_k > 0$ and most $\text{Cov}_{\text{cs}}(\beta_{i,p}, \beta_{i,k}) < 0$, the sum in (35) is negative. The empirical finding that ρ is negative for 87% of factors (Section 4) reflects this substitution structure, not coincidence.

Why $|\rho|$ increases with the Sharpe ratio. High-Sharpe-ratio factors sort assets along economically meaningful dimensions where the cross-sectional beta structure is tightly related to the omitted pricing dimensions. The cross-sectional variation in $\beta_{i,p}$ for such factors reflects genuine economic sorting, not noise, so its correlation with $\alpha_{i,p}$ is well identified and tends to be large in absolute value. Low-Sharpe-ratio factors have betas dominated by idiosyncratic variation that attenuates $|\rho|$ toward zero. The two components of the OFB therefore move in opposite directions as the Sharpe ratio increases: $\text{SD}_{\text{cs}}(\alpha)$ falls (the GRS constraint tightens) while $|\rho|$ rises (structured betas align more with the residual alphas). Section 4 documents that the absolute OFB increases with the Sharpe ratio, meaning the directional effect dominates. The wrong-sign rate nevertheless falls because $E[r_p]$ grows faster than the OFB: the wrong-sign threshold requires $|\text{OFB}| > |E[r_p]|$, which becomes harder to satisfy as $E[r_p]$ increases.

Why no observable determines the OFB. Without the APT, no-arbitrage places no constraint on $\text{Cov}_{\text{cs}}(\beta, \alpha)$ beyond the GRS bound on $\|\alpha\|$. The SDF projection (1) determines b through the moments of r_p alone (Hansen and Jagannathan 1991). The cross-sectional distribution of betas lives in a different space: it is a property of N test assets, not of the factor's time-series moments. No bridge connects the two. The Sharpe ratio constrains $\text{Var}(\varepsilon)$ through (17), the beta distribution determines $\text{Var}_{\text{cs}}(\beta)$, and neither object restricts

$\mathrm{Cov}_{\mathrm{cs}}(\beta, \alpha)$, the inner product that governs the OFB in (19).

Even under the APT (approximate factor structure), the constraint is quantitative, not qualitative. The APT bounds individual $|\alpha_i|$ by a constant ϵ that shrinks with the number of pervasive factors, but it does not constrain the cross-sectional alignment of alphas with betas. Small alphas systematically aligned with betas produce an OFB up to $\epsilon/\mathrm{SD}_{\mathrm{cs}}(\beta)$, which exceeds typical factor risk premia when $\mathrm{SD}_{\mathrm{cs}}(\beta)$ is small, as it is for most factors outside their own sorted portfolios.

Using individual assets instead of sorted portfolios does not help. As $N \to \infty$, $\mathrm{Cov}_{\mathrm{cs}}(\beta_p, \alpha)$ converges to a nonzero constant by a law of large numbers: $\sum_{k \neq p} \lambda_k \, \mathrm{Cov}_F(\beta_p, \beta_k)$, where Cov_F is the cross-firm covariance. The bias is structural, not a test-asset artifact. Firms are bundles of characteristics, so factor loadings are cross-sectionally correlated as a matter of economic organization, not of portfolio construction.

No function of observable Sharpe ratios, beta distributions, or factor correlations pins down the OFB, because no observable constrains $\rho_{\mathrm{cs}}(\beta, \alpha)$. The cross-sectional beta covariance matrix Ω_β determines the FWL amplification potential (the denominator in (41)), not the alignment term itself (the numerator). The only escape is spanning: $\varepsilon = 0$ in (1), which eliminates the alphas entirely and renders the question moot.

Corollary 3 (Worst case)**.** *Let W be any positive-definite $N \times N$ weight matrix. The weighted cross-sectional slope*

$$\gamma_1^W = E[r_p] + \frac{\tilde{\beta}' W \alpha}{\tilde{\beta}' W \tilde{\beta}}$$

satisfies $\gamma_1^W = E[r_p]$ for every finite set of test assets with $\mathrm{Var}_{\mathrm{cs}}(\beta_{i,p}) > 0$ and every positive-definite W if and only if $\mathrm{SR}(p) = \mathrm{SR}(tangency)$.

Proof. If $\mathrm{SR}(p) = \mathrm{SR}(\text{tangency})$, then $\eta_p = 0$ in the tangency projection (21), so $\alpha_{i,p} = 0$ for all assets i by (23), and $\gamma_1^W = E[r_p]$ regardless of W or the test assets. Conversely, suppose $\mathrm{SR}(p) < \mathrm{SR}(\text{tangency})$. Then $c_p > 0$ and $\eta_p \neq 0$. Consider using r_p and η_p themselves as test assets. The portfolio r_p has $(\beta_{i,p}, \beta_{i,\eta}) = (1, 0)$ and η_p has

$(\beta_{i,p}, \beta_{i,\eta}) = (0, 1)$, since $\operatorname{Cov}(r_p, \eta_p) = 0$. The alpha vector is $\alpha = (0, c_p)'$ by (23), since $\alpha_{\eta_p} = (E[r^*]/\operatorname{Var}(r^*))\operatorname{Cov}(\eta_p, \eta_p) = c_p$ and $\alpha_{r_p} = (E[r^*]/\operatorname{Var}(r^*))\operatorname{Cov}(\eta_p, r_p) = 0$. Since $\tilde{\beta} = (1/2, -1/2)'$ and $\alpha = (0, c_p)'$, we have $\tilde{\beta}'\alpha = -c_p/2 \neq 0$, and for any positive-definite W, $\tilde{\beta}'W\alpha \neq 0$. □

The corollary extends the OLS result to any weighted cross-sectional estimator, including GLS ($W = \Sigma_e^{-1}$). No weighting scheme guarantees $\gamma_1 = E[r_p]$ unless the factor spans the SDF. For any factor that does not span the SDF, there exist test-asset configurations that produce a zero or wrong-sign slope under any positive-definite weighting. In practice, the researcher does not choose test assets adversarially, but neither can the researcher verify that the cross-sectional correlation $\rho_{\text{cs}}(\beta, \alpha)$ is zero without knowing the true alphas, which requires knowing the true model.

2.12 Estimation error and the omitted factor bias

The OFB formula (5) uses population quantities: true betas $\beta_{i,p}$ and true expected returns $E[r_i]$. In practice, the researcher estimates betas from time-series regressions and uses sample averages as proxies for expected returns. This distinction raises the question of whether any empirical deviation of the Fama-MacBeth slope from $E[r_p]$ might reflect errors-in-variables (EIV) bias rather than the omitted factor bias identified above. The two biases are distinct in origin, direction, and remedy.

Suppose the researcher observes estimated betas $\hat{\beta}_{i,p} = \beta_{i,p} + u_i$, where u_i is the estimation error with $E[u_i] = 0$ and $\operatorname{Var}(u_i) = \sigma^2_{u_i}$. Denote the Fama-MacBeth estimator by $\hat{\gamma}_1^{\text{FM}}$, the time-series average of the period-by-period cross-sectional slopes. By the law of large numbers, as the number of cross-sectional periods grows, $\hat{\gamma}_1^{\text{FM}}$ converges to the cross-sectional OLS slope of $E[r_i]$ on $\hat{\beta}_{i,p}$:

$$\hat{\gamma}_1^{\text{FM}} \xrightarrow{p} \frac{\operatorname{Cov}_{\text{cs}}\left(\hat{\beta}_{i,p},\ E[r_i]\right)}{\operatorname{Var}_{\text{cs}}(\hat{\beta}_{i,p})}.$$

The estimation errors u_i are uncorrelated with both the true betas $\beta_{i,p}$ and the expected returns $E[r_i]$ in the cross-section. The condition holds exactly when u_i and $\beta_{i,p}$ are independent across assets, and approximately when betas are estimated from a long time series with homogeneous residual variance. Under this assumption, the numerator reduces to $\mathrm{Cov}_{\mathrm{cs}}(\beta_{i,p}, E[r_i])$ and the denominator to $\mathrm{Var}_{\mathrm{cs}}(\beta_{i,p}) + \bar{\sigma}_u^2$, where $\bar{\sigma}_u^2 = N^{-1}\sum_{i=1}^{N} \sigma_{u_i}^2$ is the average estimation error variance.

Remark 7. *With estimated betas, the Fama-MacBeth slope satisfies*

$$E[\hat{\gamma}_1^{\mathrm{FM}}] \approx \lambda\, \gamma_1, \tag{36}$$

where γ_1 *is the population slope from* (5) *and*

$$\lambda = \frac{\mathrm{Var}_{\mathrm{cs}}(\beta_{i,p})}{\mathrm{Var}_{\mathrm{cs}}(\beta_{i,p}) + \bar{\sigma}_u^2} \in (0,1]$$

is the errors-in-variables attenuation factor. The estimated slope decomposes as

$$E[\hat{\gamma}_1^{\mathrm{FM}}] - E[r_p] \approx -\underbrace{(1-\lambda)\, E[r_p]}_{\text{EIV attenuation}} - \underbrace{\lambda \cdot \mathrm{OFB}}_{\text{omitted factor bias}}, \tag{37}$$

where $\mathrm{OFB} = \mathrm{Cov}_{\mathrm{cs}}(\beta_{i,p},\ \mathrm{Cov}(\varepsilon, r_i)/E[m]) \,/\, \mathrm{Var}_{\mathrm{cs}}(\beta_{i,p})$.

The two terms are distinct in both direction and origin. The EIV attenuation $(1 - \lambda)\, E[r_p]$ always pushes the estimated slope toward zero. EIV attenuation is a finite-sample phenomenon that vanishes as T grows ($\bar{\sigma}_u^2 \to 0$ and $\lambda \to 1$) or when betas are estimated precisely, as with sorted portfolios. The OFB can be positive, negative, or zero. It is a property of the population and persists regardless of sample size.

For 25 sorted portfolios as test assets with $T = 600$ months of data, the cross-sectional dispersion in market betas is typically of order 10^{-1}, while the average beta estimation error variance is of order 10^{-3}, so λ is close to one and EIV attenuation is negligible. For individual stocks estimated over short windows ($T = 60$ months), $\bar{\sigma}_u^2$ can be an order of magnitude larger, making λ substantially below one and both biases empirically relevant. Section 4 quantifies these magnitudes for a range of test assets.

The Shanken (1992) correction adjusts the standard errors of the Fama-MacBeth estimator to account for the estimation error in betas; it does not adjust the point estimate, and it does not remove the OFB. The Fama-MacBeth point estimate still converges to $\gamma_1 = E[r_p] + \text{OFB}$, not to $E[r_p]$. A researcher who applies the Shanken correction and reports a significant risk premium is testing whether $\gamma_1 \neq 0$, which is a joint hypothesis about $E[r_p]$ and the OFB, not a test of whether betas on p command a cross-sectional price equal to $E[r_p]$.

The decomposition (37) also clarifies the role of test asset choice in a second way. Beyond determining the OFB, the choice of test assets affects the attenuation factor λ through $\text{Var}_{\text{cs}}(\beta_{i,p})$: a set of test assets with greater cross-sectional dispersion in betas produces a higher λ and less EIV attenuation. Researchers thus face a tension. Using a small number of well-diversified portfolios minimizes $\bar{\sigma}_u^2$ but may also reduce $\text{Var}_{\text{cs}}(\beta_{i,p})$. Using individual stocks maximizes the cross-sectional variation in betas but increases $\bar{\sigma}_u^2$. Neither choice addresses the OFB, which depends on the cross-sectional distribution of $\text{Cov}(\varepsilon, r_i)$; longer time series do not reduce it.

2.13 Extension to multi-factor models

Asset pricing models use multiple factors simultaneously. The OFB characterization extends to the multi-factor setting, with the Frisch-Waugh-Lovell (FWL) theorem governing the alignment term for each individual factor.

Consider K factors $f = (f_1, \ldots, f_K)'$, each an excess return. Project the SDF onto the span of f:

$$m = a + b'f + \varepsilon, \qquad \text{Cov}(\varepsilon, f) = 0, \tag{38}$$

where $b = -E[m]\, \Sigma_f^{-1}\, E[f]$ and $\Sigma_f = \text{Var}(f)$. The residual ε now captures the SDF variation unspanned by all K factors jointly.

Substituting (38) into the pricing condition $E[m \cdot r_i] = 0$ and dividing by $-E[m]$ yields the multi-factor pricing equation:

$$E[r_i] = \beta_i' \, E[f] - \frac{\operatorname{Cov}(\varepsilon, r_i)}{E[m]}, \tag{39}$$

where $\beta_i = \Sigma_f^{-1} \operatorname{Cov}(f, r_i)$ is the $K \times 1$ vector of multiple-regression betas. The pricing error of asset i with respect to the K-factor model is $\alpha_i = E[r_i] - \beta_i' \, E[f] = -\operatorname{Cov}(\varepsilon, r_i)/E[m]$, exactly as in the single-factor case but with ε defined relative to the full factor set.

A cross-sectional regression of $E[r_i]$ on β_i yields the $K \times 1$ vector of Fama-MacBeth slopes.

Proposition 2. *The population Fama-MacBeth slope vector in the K-factor model satisfies*

$$\gamma = E[f] + (\tilde{B}'\tilde{B})^{-1}\tilde{B}'\alpha, \tag{40}$$

where $\tilde{B}$ is the $N \times K$ matrix of demeaned betas ($\tilde{\beta}_{i,k} = \beta_{i,k} - \bar{\beta}_k$) and $\alpha = (\alpha_1, \ldots, \alpha_N)'$ is the vector of pricing errors. The alignment term for factor k satisfies

$$\gamma_k - E[f_k] = \frac{\sum_{i=1}^{N} \tilde{\beta}_{i,k}^{\perp} \, \alpha_i}{\sum_{i=1}^{N} (\tilde{\beta}_{i,k}^{\perp})^2}, \tag{41}$$

where $\tilde{\beta}_{i,k}^{\perp}$ is the residual from the cross-sectional regression of $\tilde{\beta}_{i,k}$ on all other demeaned betas $\tilde{\beta}_{i,-k}$.

Proof. Equation (40) follows from the cross-sectional OLS formula. Writing $E[r_i] = \beta_i' \, E[f] + \alpha_i$ and substituting into the multivariate OLS expression $\gamma = (\tilde{B}'\tilde{B})^{-1}\tilde{B}' \, E[r]$:

$$\gamma = (\tilde{B}'\tilde{B})^{-1}\tilde{B}'\big(\tilde{B} \, E[f] + \bar{\beta}' \, E[f] \, \mathbf{1} + \alpha\big) = E[f] + (\tilde{B}'\tilde{B})^{-1}\tilde{B}'\alpha,$$

where the $\bar{\beta}' \, E[f] \, \mathbf{1}$ term vanishes because $\tilde{B}'\mathbf{1} = 0$ by construction. Equation (41) then follows from the Frisch-Waugh-Lovell theorem: the k-th coefficient in a multivariate OLS regression equals the coefficient from regressing the dependent variable on the residual from projecting the k-th regressor onto all others. $\square$

The FWL decomposition (41) reveals a structural difference from the single-factor case. The OFB for factor k depends not on the raw demeaned beta $\tilde{\beta}_{i,k}$ but on the partialled-out component $\tilde{\beta}_{i,k}^{\perp}$: the cross-sectional variation in factor k's beta that is not explained by the betas on all other factors. When factors are correlated, their betas are correlated across assets, so $\text{Var}_{\text{cs}}(\tilde{\beta}_{i,k}^{\perp})$ can be much smaller than $\text{Var}_{\text{cs}}(\tilde{\beta}_{i,k})$. This reduction in variance has two consequences. The denominator in (41) shrinks, amplifying any nonzero numerator and producing large alignment terms from even modest pricing errors. The numerator is driven by the covariance between the pricing errors and the unique component of factor k's beta, so only the part of the alpha vector that aligns with this residual beta matters for the alignment term.

The Sharpe ratio decomposition generalizes directly:

$$\text{SR}^2(\text{tangency}) = E[f]'\Sigma_f^{-1}E[f] + \frac{\text{Var}(\varepsilon)}{E[m]^2}, \tag{42}$$

where $E[f]'\Sigma_f^{-1}E[f] = \text{SR}^2(f)$ is the squared Sharpe ratio of the optimal combination of the K factors. As in the single-factor case, the tangency portfolio decomposes into a component spanned by f and an orthogonal residual η, with $\varepsilon = b^*\eta$. The Cauchy-Schwarz bound on each $|\,\text{Cov}(\varepsilon, r_i)|$ constrains individual alphas through $\text{Var}(\varepsilon)$, but the OFB per factor depends on how those alphas project onto the FWL residual betas $\tilde{\beta}_{i,k}^{\perp}$, a quantity that $\text{SR}(f)$ does not constrain.

This extension connects to the discussion in Kandel and Stambaugh (1995), who noted in passing that their repackaging result generalizes to multi-factor models. The present result differs in focus: Kandel and Stambaugh showed that changing the test assets can produce arbitrary slopes, whereas (41) characterizes the OFB for a fixed cross-section. Multicollinearity among factor betas amplifies the OFB for individual factors through the FWL denominator, even when the aggregate pricing errors $\alpha'\Sigma_e^{-1}\alpha$ are small. Adding a $(K+1)$-th factor reduces $\mathrm{Var}(\varepsilon)$ by expanding the spanned subspace, but simultaneously introduces additional collinearity among the betas, potentially shrinking $\mathrm{Var}_{\mathrm{cs}}(\tilde{\beta}^{\perp}_{i,k})$ for the original factors. The net effect on the OFB is ambiguous. Section 4 documents these patterns empirically.

Remark 8 (Nontradable factors)**.** *The preceding results assume r_p is an excess return, so the pricing condition $E[m \cdot r_p] = 0$ pins down $b = -E[m]\,E[r_p]/\,\mathrm{Var}(r_p)$. For a nontradable factor f (not a return), $E[m \cdot f] = 0$ need not hold, and the projection coefficient $b = \mathrm{Cov}(m, f)/\,\mathrm{Var}(f)$ is a primitive of the SDF, not determined by the factor's moments alone. Define the population price of risk*

$$\lambda_f \;=\; -\frac{\mathrm{Cov}(m, f)}{E[m]} \;=\; -\frac{b\;\mathrm{Var}(f)}{E[m]}. \tag{43}$$

Substituting the SDF projection $m = a + b\,f + \varepsilon$ into $E[m \cdot r_i] = 0$ and dividing by $-E[m]$ gives the pricing equation

$$E[r_i] \;=\; \beta_{i,f}\,\lambda_f \;-\; \frac{\mathrm{Cov}(\varepsilon, r_i)}{E[m]}, \tag{44}$$

where $\beta_{i,f} = \mathrm{Cov}(f, r_i)/\,\mathrm{Var}(f)$. This is identical to (4) *with λ_f replacing $E[r_p]$. The Fama-MacBeth slope inherits the same OFB structure:*

$$\gamma_1 \;=\; \lambda_f \;+\; \frac{\mathrm{Cov}_{\mathrm{cs}}(\beta_{i,f},\; \alpha_{i,f})}{\mathrm{Var}_{\mathrm{cs}}(\beta_{i,f})}, \tag{45}$$

where $\alpha_{i,f} \;=\; E[r_i] - \beta_{i,f}\,\lambda_f \;=\; -\,\mathrm{Cov}(\varepsilon, r_i)/E[m]$. The decomposition into magnitude and direction (Corollary 1), the Cauchy-Schwarz bound (Proposition 1), and the universality result (Corollary 3) all carry over with λ_f in place of $E[r_p]$.

The difference is observability. For a tradable factor, $\lambda_f = E[r_p]$ is estimable from the factor's own time series. For a nontradable factor, λ_f depends on $\mathrm{Cov}(m, f)$, which requires knowledge of the SDF. The mimicking portfolio $r_f^* \;=\; \mathrm{Cov}(f, R)'\,\mathrm{Var}(R)^{-1}R$, where R is *the vector of asset returns, satisfies* $E[r_f^*] = \lambda_f$ *and recovers the price of risk, but does not eliminate the OFB: the residual* ε *in* (44) *captures the SDF variation unspanned by* f*, and the mimicking portfolio inherits this residual because it spans only what* f *spans. The tradable factor risk premia (TFRP) approach of Bryzgalova et al. (2023) constructs* $\lambda^* = \mathrm{Cov}(f, R)\ \mathrm{Var}(R)^{-1}\, E[R]$*, which equals* $E[r_p]$ *for tradable factors and* $E[r_f^*]$ *for nontradable ones. This bypasses the Fama-MacBeth regression but does not bypass the OFB, because the minimum-variance SDF implied by the factors,* $m^* = 1 - \gamma^{*\prime}(R - \mu_R)$*, differs from the true SDF by the unspanned component* ε^**, which drives the pricing errors* $\alpha_{i,f}$ *and hence the cross-sectional correlation between betas and alphas. A more direct route to the SDF, the tangency portfolio with weights* $\gamma^* = \mathrm{Var}(R)^{-1}E[R]$*, requires knowledge of population expected returns* $E[R]$*, the very object cross-sectional methods aim to estimate. Ex post, the sample covariance matrix* $\hat{V}_R$ *is singular whenever the number of test assets* N *exceeds the sample size* T*, so the sample tangency portfolio does not exist without regularization. These constraints force practitioners toward small test-asset sets (*$N \ll T$*), effectively projecting the SDF onto a low-dimensional subspace, which is itself a factor model. The OFB is a property of the factor's incompleteness relative to the SDF, not of the estimation method.*

3 Monte Carlo Validation

The preceding results are exact population statements, not asymptotic approximations. A Monte Carlo exercise serves two purposes. First, it confirms the formula numerically in a fully specified economy where every quantity is known. Second, it shows the predictions visually: that the Sharpe ratio is orthogonal to the OFB, that test-asset choice drives the bias, and that FWL amplification can reach a factor of five in multi-factor models. The design uses independent factors and homogeneous idiosyncratic variance, so the parameters determine every analytical prediction exactly. Section 4 confirms the same patterns in real data with all the complications of empirical factor structures.

3.1 Design

Three independent traded factors generate all expected returns:

$$r_{i,t} = \beta_{i,1}\, f_{1,t} + \beta_{i,2}\, f_{2,t} + \beta_{i,3}\, f_{3,t} + e_{i,t}, \qquad i = 1, \ldots, N, \tag{46}$$

where $f_k \sim N(\mu_k, \sigma_k^2)$ independently across factors, $e_{i,t} \sim N(0, \sigma_e^2)$ independently across assets and time, and $N = 100$. The true SDF is the minimum-variance SDF in this three-factor economy, so all assets are priced exactly: $E[r_i] = \beta_{i,1}\, \mu_1 + \beta_{i,2}\, \mu_2 + \beta_{i,3}\, \mu_3$ and $\alpha_i = 0$ relative to the true model.

The researcher observes only f_1 and estimates a single-factor model, omitting f_2 and f_3. Since the factors are independent, the time-series regression recovers the beta on f_1 without bias, but the pricing error relative to the misspecified model is

$$\alpha_i = \beta_{i,2}\, \mu_2 + \beta_{i,3}\, \mu_3, \tag{47}$$

which is nonzero whenever $\beta_{i,2}$ or $\beta_{i,3}$ are not cross-sectionally constant.

I draw the beta vectors $(\beta_{i,1}, \beta_{i,2}, \beta_{i,3})$ once from a trivariate normal distribution with mean $(1, 1, 1)$ and a covariance matrix parameterized by the cross-sectional correlations ρ_{12}, ρ_{13}, and ρ_{23} between betas on different factors. These correlations are the experimental levers: ρ_{12} determines how the omitted factor f_2's pricing contribution distributes across the included factor's beta spectrum.

Substituting (47) into the population OFB formula (19) gives the population Fama-MacBeth slope under single-factor misspecification:

$$\gamma_1 = \mu_1 + \mu_2\, \rho_{12}\, \frac{s_2}{s_1} + \mu_3\, \rho_{13}\, \frac{s_3}{s_1}, \tag{48}$$

where s_k denotes the cross-sectional standard deviation of $\beta_{i,k}$. The OFB depends on the omitted factors' risk premia (μ_2, μ_3) and the cross-sectional beta correlations (ρ_{12}, ρ_{13}), but is independent of μ_1 and σ_1, and therefore of $\mathrm{SR}(f_1) = \mu_1/\sigma_1$.

Table 1 reports the baseline parameter values. All returns are in monthly percentage points. The annualized factor Sharpe ratios are modest (0.26–0.35), and the tangency Sharpe ratio of 0.53 is conservative relative to empirical estimates. I set $\rho_{13} = \rho_{23} = 0$ at baseline, so the OFB reduces to $\mu_2\,\rho_{12}\,(s_2/s_1)$. All analytical predictions use the realized cross-sectional moments of the drawn betas, not the distributional moments, to eliminate any finite-N discrepancy.

Table 1: Monte Carlo parameters

Parameter	Value	Implied quantity
μ_1, μ_2, μ_3	0.50, 0.40, 0.30	$\mathrm{SR}_1 = 0.35$, $\mathrm{SR}_2 = 0.31$, $\mathrm{SR}_3 = 0.26$ (ann.)
$\sigma_1, \sigma_2, \sigma_3$	5.0, 4.5, 4.0	SR(tangency) = 0.53 (ann.)
σ_e	3.0	EIV attenuation $\bar{\lambda} \approx 0.97$ at $T = 600$
$s_1 = s_2 = s_3$	0.3	Cross-sectional SD of $\beta_{i,k}$
ρ_{12}	0.5 (baseline)	OFB ≈ 0.15 (realized betas)
N	100	$T = 600$, $S = 1{,}000$ replications

3.2 Results

Experiment 1: Varying the Sharpe ratio at constant OFB. I vary μ_1 from 0.10 to 1.50 (% per month), changing the annualized Sharpe ratio of f_1 from 0.07 to 1.04, while holding all other parameters at their baseline values. Since the OFB does not depend on μ_1, the bias $\gamma_1 - \mu_1$ is constant at 0.15 across the entire range.

Figure 1 Panel A plots the bias ($\gamma_1 - E[f_1]$) against the annualized Sharpe ratio. The analytical prediction (dashed line) is a flat line at 0.15. The simulated Fama-MacBeth slopes (averaged over 1,000 replications with $T = 600$ and population betas) match the prediction to within sampling variation. The Sharpe ratio varies by a factor of fifteen; the bias does not move. Repeating the exercise with betas estimated from the simulated time series ($T = 600$) produces slopes attenuated by the factor $\bar{\lambda} \approx 0.96$, consistent with Remark 7.

Experiment 2: Varying the OFB at constant Sharpe ratio. I fix $\mu_1 = 0.50$ and $\sigma_1 = 5.0$ (so $\text{SR}(f_1) = 0.35$ annualized) and vary ρ_{12} from -0.9 to $+0.9$, redrawing betas for each value. Varying ρ_{12} changes the OFB while the Sharpe ratio does not move.

Panel B plots the FM slope against ρ_{12}. The slope crosses through $E[f_1] = 0.50$ near $\rho_{12} = 0$ and swings in both directions as the beta correlation moves away from zero. The simulated slopes match the analytical prediction at each point. At the same Sharpe ratio, the FM regression can overestimate, correctly estimate, or underestimate the factor's risk premium, depending on how the omitted factor's betas correlate with the included factor's betas across the cross-section.

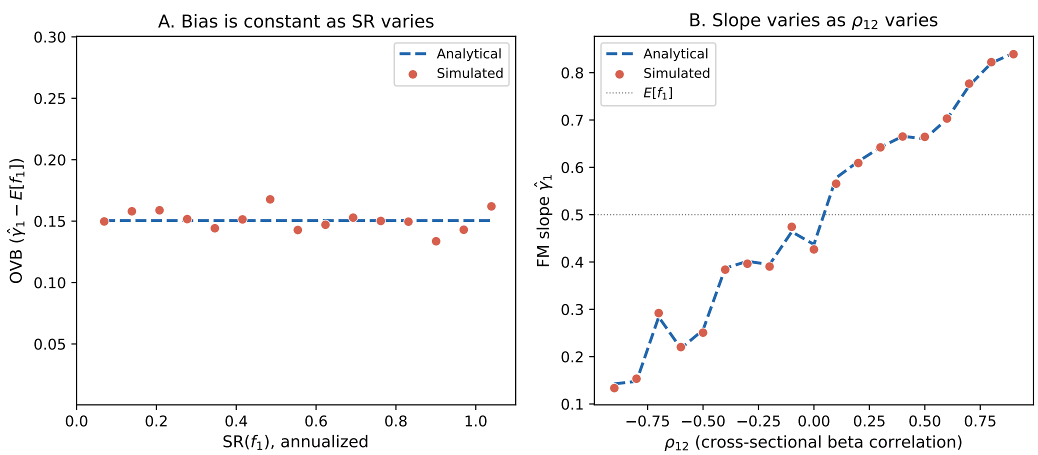

Figure 1: Monte Carlo validation. Panel A: the OFB ($\gamma_1 - E[f_1]$) is constant at 0.15 as the Sharpe ratio varies from 0.07 to 1.04. Panel B: the FM slope varies with ρ_{12}, crossing through $E[f_1] = 0.50$ at $\rho_{12} = 0$, while the Sharpe ratio is fixed at 0.35. Dots are simulated FM slopes (1,000 replications, $T = 600$, population betas); dashed lines are analytical predictions from equation (48). Three-factor DGP with $N = 100$ test assets; see Table 1 for parameters.

Experiment 3: Test-asset dependence. To isolate the test-asset mechanism, I replace the random betas with deterministic grids. I construct four test-asset configurations of $N = 50$ assets each, using the same factor f_1 (same Sharpe ratio, same expected return) but different cross-sectional structures for $\beta_{i,2}$. In all configurations, $\beta_{i,1}$ is equally spaced from 0.4 to 1.6 and $\beta_{i,3} = 1.0$ for all i (constant, contributing no OFB). The deterministic construction ensures exact analytical predictions; the empirical illustration in Section 4 confirms the same pattern with observed betas and real test assets.

- *Orthogonal*: $\beta_{i,2} = 1.0$ for all i. The omitted factor's beta is constant, so $\mathrm{Cov}_{\mathrm{cs}}(\beta_{i,1}, \alpha_i) = 0$ and $\gamma_1 = E[f_1]$.
- *Aligned*: $\beta_{i,2} = \beta_{i,1}$. The omitted factor's betas are perfectly positively correlated with the included factor's betas, so $\gamma_1 = E[f_1] + \mu_2 = 0.90$.
- *Opposed*: $\beta_{i,2} = 2.0 - \beta_{i,1}$. The correlation is -1, so $\gamma_1 = E[f_1] - \mu_2 = 0.10$.
- *Independent*: $\beta_{i,2}$ drawn independently from $U(0.4, 1.6)$. The cross-sectional correlation is approximately zero, so $\gamma_1 \approx E[f_1]$.

Figure 2 displays the results. The same factor, with the same Sharpe ratio and expected return, produces population FM slopes ranging from 0.10 to 0.90 depending on the test assets. The analytical predictions match the cross-sectional regressions. The FM slope is a joint property of the factor and the test assets, not of the factor alone.

Experiment 4: Multi-factor FWL amplification. I include both f_1 and f_2 as factors, omitting only f_3, and set $\rho_{13} = 0.4$ (so $\beta_{i,3}$ is correlated with $\beta_{i,1}$, creating nonzero OFB for factor 1 even in the two-factor model) and $\rho_{23} = 0$. I vary ρ_{12}, the cross-sectional correlation between $\beta_{i,1}$ and $\beta_{i,2}$, from 0 to 0.9.

From the FWL structure (Proposition 2), the OFB for factor 1 in the two-factor model is, with $\rho_{23} = 0$,

$$\gamma_1 - \mu_1 = \frac{\mu_3\, s_3\, \rho_{13}}{s_1\,(1 - \rho_{12}^2)}, \tag{49}$$

where the denominator $1 - \rho_{12}^2$ is the FWL deflation of the cross-sectional variance. As ρ_{12} increases, the partialled-out betas $\tilde{\beta}_{i,1}^{\perp}$ lose cross-sectional variation, amplifying the bias per unit of mispricing by the factor $1/(1 - \rho_{12}^2)$, while the model's aggregate squared Sharpe

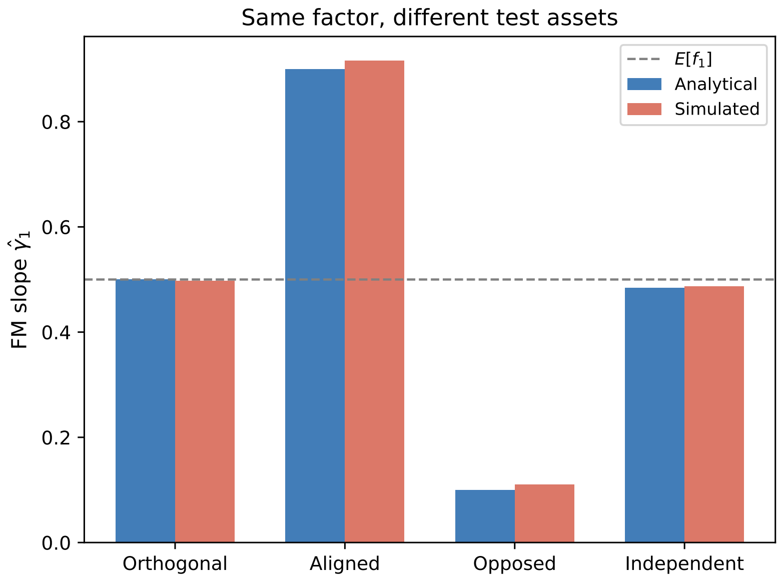

Figure 2: Test-asset dependence: population FM slope for the same factor ($E[f_1] = 0.50$, SR $= 0.35$) across four test-asset configurations that differ only in how $\beta_{i,2}$ (the omitted factor's beta) relates to $\beta_{i,1}$ (the included factor's beta). The dashed line marks $E[f_1]$. All analytical predictions (bars) match the cross-sectional regressions exactly.

ratio $\mathrm{SR}^2(f_1, f_2) = \mu_1^2/\sigma_1^2 + \mu_2^2/\sigma_2^2$ remains constant (since it depends on factor means and variances, not on cross-sectional beta correlations).

Table 2 reports the results. At $\rho_{12} = 0$, the OFB for factor 1 is 0.07. At $\rho_{12} = 0.9$, the FWL amplification factor reaches 5.3 and the OFB grows to 0.59, exceeding the factor's own risk premium. The model's aggregate Sharpe ratio is unchanged throughout. The population GLS slope shows OFB of similar magnitude to the OLS slope in every row, confirming that the bias originates in the misspecification, not in the estimator's weighting scheme. Adding factors to improve the model's Sharpe ratio can worsen the bias for individual prices of risk when the factors' betas are correlated across the cross-section.

4 Empirical Illustration

The OFB formula holds for any portfolio with a nonzero expected return. This section measures both components of γ_1 across the population of candidate factors to characterize

Table 2: FWL amplification in a two-factor model

ρ_{12}	FWL ratio $1/(1-\rho_{12}^2)$	Analytical OFB $\gamma_1 - \mu_1$	Simulated OFB $\hat{\gamma}_1 - \mu_1$	GLS OFB $\gamma_1^{\text{GLS}} - \mu_1$	Model SR (ann.)
0.0	1.00	0.07	0.06	0.08	0.46
0.3	1.10	0.15	0.16	0.15	0.46
0.5	1.33	0.19	0.20	0.20	0.46
0.7	1.96	0.23	0.21	0.25	0.46
0.8	2.78	0.28	0.29	0.26	0.46
0.9	5.26	0.59	0.62	0.61	0.46

Notes. Two-factor model using f_1 and f_2, omitting f_3. The cross-sectional correlation between $\beta_{i,1}$ and $\beta_{i,3}$ is fixed at $\rho_{13} = 0.4$; $\rho_{23} = 0$. The FWL ratio $\text{Var}_{\text{cs}}(\beta_{i,1})/\text{Var}_{\text{cs}}(\tilde{\beta}_{i,1}^{\perp})$ measures the amplification of bias from multicollinearity. The model's aggregate Sharpe ratio does not depend on ρ_{12}. Analytical OFB is the population OFB computed from the realized cross-sectional moments of the drawn betas. $N = 100$, $T = 600$, $S = 1{,}000$ replications.

how the Sharpe ratio relates to cross-sectional pricing empirically, not to evaluate single-factor models as practical pricing specifications. The empirical sample is 212 long-short factor portfolios from the Chen and Zimmermann (2022) Open Source Asset Pricing database. For each factor, betas of test assets come from time-series regressions, followed by a cross-sectional regression of average excess returns on these betas. If the Sharpe ratio determined γ_1, the estimated slope $\hat{\gamma}_1$ should track $E[r_p]$ across factors, so that factors with high Sharpe ratios have slopes close to their own expected returns. The data reject both predictions. Section 4.5 extends the analysis to multi-factor models.

The test assets are all sorted portfolios available in the database: for each of the 212 signals, I include every non-long-short sorted portfolio (deciles, quintiles, or terciles, depending on the signal), yielding 1,277 portfolios in total. Using the full set of sorted portfolios avoids arbitrary choices about which characteristics to include. I estimate betas four ways: full-sample OLS, rolling 60-month windows, daily returns over 252-day windows, and daily returns over 1260-day windows. Rolling and daily betas are lagged by one month to avoid look-ahead bias. All returns are monthly and expressed in percentage points. I require a minimum of 120 months of non-missing returns for each factor, yielding 212 factors.

4.1 Sharpe Ratios and Fama-MacBeth Slopes

Figure 3 plots the absolute Sharpe ratio against the estimated Fama-MacBeth slope for each of the 212 factors, using full-sample betas and all 1,277 sorted portfolios as test assets. The pattern is a flat cloud; the correlation between $|SR|$ and $\hat{\gamma}_1$ ranges from 0.10 to 0.15 across beta estimation methods (Table 3). Factors with annualized Sharpe ratios above 1.5 produce FM slopes that range from negative to positive, with no tendency toward their own expected returns.

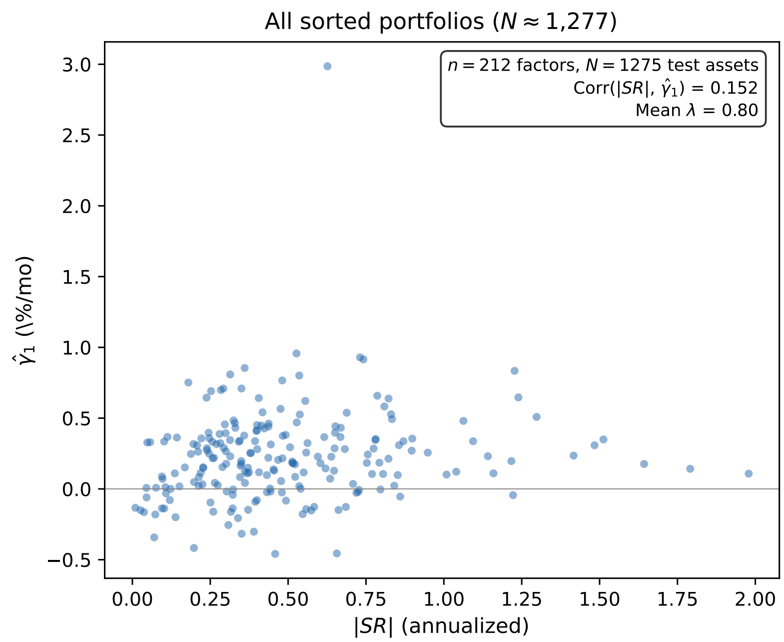

Figure 3: Absolute Sharpe ratio versus Fama-MacBeth slope for 212 long-short factors, using full-sample betas and all 1,277 sorted portfolios as test assets. The scatter shows no systematic relationship between $|SR|$ and $\hat{\gamma}_1$.

Table 3 reports summary statistics across four beta estimation methods, all using the same 1,277 test assets.

Between 22% and 31% of factors have wrong-sign FM slopes, meaning $\hat{\gamma}_1$ and the factor's average excess return $\bar{r}_p$ have opposite signs. A factor with a positive expected return produces a negative estimated price of risk, or vice versa. This finding is consistent with Remark 2: wrong-sign slopes are possible for any factor that does not span the SDF, and

Table 3: Fama-MacBeth slope vs. factor risk premium

Estimator	N	$\bar{\lambda}$	Corr($\lvert SR \rvert$, $\hat{\gamma}_1$)	Wrong sign	Med. $\hat{\gamma}_1/\bar{r}_p$
Panel A: Period-by-period FM slopes					
Full sample	1275	0.80	0.152	22.6%	0.39
Rolling 60mo	1277	0.56	0.123	30.7%	0.27
Daily 252d	1277	0.75	0.112	21.7%	0.54
Daily 1260d	1277	0.88	0.103	24.1%	0.65
Panel B: Population cross-sectional slopes					
Population OLS	1275	–	0.185	21.2%	0.32
Population GLS	1275	–	0.180	23.1%	0.22

Notes. Each row reports cross-factor summary statistics from Fama-MacBeth regressions of 212 long-short factors from Chen and Zimmermann (2022) against all 1,277 sorted portfolios. Panel A reports the time-series average of period-by-period cross-sectional slope estimates $\hat{\gamma}_1$. Panel B reports population slopes from a single cross-sectional regression of mean excess returns on full-sample betas: OLS for the standard estimator and GLS with diagonal $\hat{\Sigma}_e$ (residual variances from the time-series regressions) following Kandel and Stambaugh (1995). $\bar{\lambda}$ is the mean EIV attenuation factor $\mathrm{Var}_{\mathrm{cs}}(\hat{\beta})/[\mathrm{Var}_{\mathrm{cs}}(\hat{\beta}) + \bar{\sigma}_u^2]$. "Wrong sign" reports the fraction of factors where $\hat{\gamma}_1$ and $\bar{r}_p$ have opposite signs. Rolling and daily betas are lagged by one month.

their prevalence does not diminish as $|SR|$ increases.

To benchmark these wrong-sign rates, I simulate the null of correct specification for each factor. Under the null, the true model is the single factor with $\alpha_i = 0$ for all test assets, calibrated to each factor's actual moments (expected return, variance, beta dispersion, residual variance, sample length). Across 1,000 replications per factor, the average wrong-sign rate under the null is 4.9%, concentrated almost entirely in the lowest $|SR|$ quintile (20.5%) where weak signals and short samples allow finite-sample noise to flip the sign. For factors with $|SR| > 0.35$, the null wrong-sign rate is below 1%. The observed wrong-sign rate of 22.6% exceeds the null by nearly 18 percentage points: finite-sample noise accounts for roughly one-fifth of the observed wrong-sign problem, and the remainder requires misspecification.

4.2 Errors-in-Variables versus Omitted Factor Bias

The weak relationship between Sharpe ratios and FM slopes might reflect errors-in-variables attenuation from noisy beta estimates. Remark 7 provides the framework to evaluate this objection. The total deviation of the FM slope from $E[r_p]$ decomposes into an EIV attenuation term and an OFB term, scaled by the attenuation factor λ ((37)).

The attenuation factor λ varies across beta estimation methods, as reported in Table 3. Moving from rolling 60-month betas to daily 1260-day betas raises $\bar{\lambda}$ from 0.56 to 0.88. If EIV were the only source of bias, reducing attenuation should bring the median ratio $\hat{\gamma}_1/\overline{r}_p$ close to one and eliminate wrong-sign slopes (EIV attenuates toward zero but cannot flip signs). The data contradict this prediction: even with daily 1260-day betas ($\bar{\lambda} = 0.88$), the median ratio $\hat{\gamma}_1/\overline{r}_p$ is 0.65, and 24% of factors retain wrong-sign slopes. The gap between the predicted EIV-only ratio and the observed ratio reflects the OFB component. Pure EIV attenuation cannot produce wrong signs: multiplying a positive $E[r_p]$ by $\lambda \in (0, 1]$ cannot produce a negative slope. Wrong signs require the OFB to exceed $E[r_p]$ in magnitude, pushing the slope past zero and into the opposite-sign region.

The pattern across beta methods reinforces this conclusion. Moving from rolling 60-month betas to daily 1260-day betas reduces $1 - \lambda$ from 0.44 to 0.12, but wrong-sign rates fall only from 31% to 24%. The OFB drives the disconnect between Sharpe ratios and estimated prices of risk: wrong-sign rates fall only from 31% to 24% as the attenuation factor rises from 0.56 to 0.88.

Figure 4 isolates the OFB directly. For each factor, I compute the full-sample cross-sectional slope: a single regression of full-sample mean excess returns on full-sample betas, using all 1,277 sorted portfolios as test assets. Panel A's "full sample" row and Panel B's "Population OLS" row both use full-sample betas, but they differ in one respect. Panel A follows the standard Fama-MacBeth procedure: it runs a cross-sectional regression each month using that month's realized returns and the same full-sample betas, then averages the monthly slopes. Each monthly slope uses the same $\hat{\beta}_i$ against different realized $r_{i,t}$, producing month-to-month variation in $\hat{\gamma}_{1,t}$; their average converges to the ratio $\operatorname{Cov}_{\text{cs}}(\hat{\beta}, \bar{r})/\operatorname{Var}_{\text{cs}}(\hat{\beta})$, which inherits EIV attenuation ($\lambda = 0.80$) because $\operatorname{Var}_{\text{cs}}(\hat{\beta})$ includes estimation error variance. Panel B computes the population OLS slope directly from full-sample moments, treating the betas as population quantities ($\lambda = 1$ by construction), so any deviation of γ_1 from $E[r_p]$ is entirely OFB. The wrong-sign rates are close (22.6% versus 21.2%) because the 20% EIV attenuation primarily affects factors near zero, occasionally pulling a barely positive slope to negative or vice versa, without changing the overall pattern. The figure plots the full-sample γ_1 against $E[r_p]$, with points colored by Sharpe-ratio quintile; under correct specification, all points would lie on the 45-degree line. Instead, the cloud deviates from the line: 22.6% of factors have wrong-sign slopes, matching the rate in Table 3 for full-sample betas and confirming that the OFB accounts for the deviation. The color gradient reveals that high-Sharpe-ratio factors (dark blue) are no closer to the 45-degree line than low-Sharpe-ratio factors (light blue): the OFB does not shrink systematically with the Sharpe ratio. Because full-sample betas set $\lambda = 1$, the OFB formula (9) is an identity: $\gamma_1^{\text{pop}} - E[r_p] = \operatorname{Cov}_{\text{cs}}(\beta, \alpha)/\operatorname{Var}_{\text{cs}}(\beta)$ holds exactly for each factor, and the predicted OFB matches the observed deviation with $R^2 = 1$.

4.3 Anatomy of the Omitted Factor Bias

Corollary 1 decomposes the OFB into a magnitude term, $\text{SD}_{\text{cs}}(\alpha)/\text{SD}_{\text{cs}}(\beta)$, that the Sharpe ratio constrains, and a directional term, $\rho_{\text{cs}}(\beta_{i,p}, \alpha_{i,p})$, that it does not. This subsection measures the two components across the 212 factors.

For each factor, I compute full-sample betas and alphas from time-series regressions of the 1,277 sorted portfolios on the factor, then compute $\hat{\rho} = \operatorname{Corr}_{\text{cs}}(\hat{\beta}_{i,p}, \hat{\alpha}_{i,p})$.

The directional term is predominantly negative. Across the 212 factors, $\hat{\rho}$ is negative for 185 factors (87%). The mean is -0.22 (median -0.22, standard deviation 0.21), with a range from -0.78 to $+0.35$. For 87% of single-factor models, assets with high factor exposure earn less than their betas predict. This finding generalizes the betting-against-beta pattern from

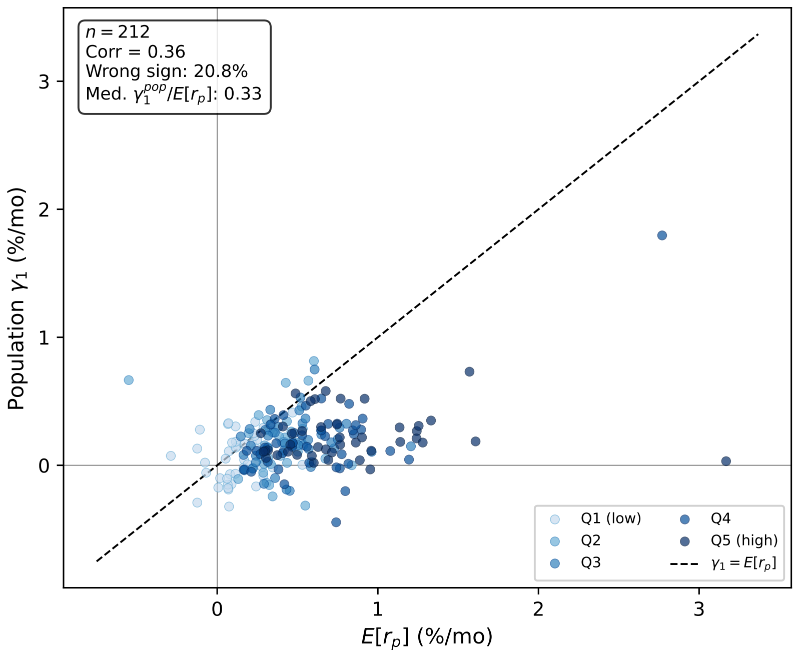

Figure 4: Full-sample cross-sectional slope versus factor expected return for 212 long-short factors, colored by $|SR(p)|$ quintile (light to dark blue). The slope is computed from a single regression of full-sample mean excess returns on full-sample betas using all 1,277 sorted portfolios. By construction, $\lambda = 1$ (no errors-in-variables). The dashed line marks $\gamma_1 = E[r_p]$. Deviations from the line are entirely omitted factor bias. High-Sharpe-ratio factors are no closer to the line than low-Sharpe-ratio factors.

the CAPM to the entire factor zoo.

High Sharpe ratios predict more directional misalignment. The theory allows ρ to take any value regardless of the Sharpe ratio. Practitioners might hope that high-Sharpe-ratio factors produce smaller $|\rho|$. The data reject this hope: $\mathrm{Corr}(|\hat{\rho}|, |SR|) = 0.49$ across the 212 factors. Table 4 reports averages by $|SR|$ quintile. Factors in the lowest quintile have a mean $|\hat{\rho}|$ of 0.15; those in the highest have 0.37. The correlation reflects a selection effect: factors with high Sharpe ratios tend to be constructed from characteristics that sort stocks along dimensions correlated with other priced risks, generating directional misalignment.

Table 4: Cross-sectional correlation $|\hat{\rho}|$ by Sharpe ratio quintile

$\|SR\|$ quintile	Mean $\|SR\|$	Mean $\|\hat{\rho}\|$	Median $\hat{\rho}$	Mean $\|\text{OFB}\|$
Q1 (low)	0.14	0.15	−0.06	0.17
Q2	0.30	0.20	−0.20	0.30
Q3	0.42	0.21	−0.19	0.31
Q4	0.60	0.30	−0.31	0.44
Q5 (high)	1.02	0.37	−0.38	0.60

Notes. Each quintile contains approximately 42 factors sorted by absolute annualized Sharpe ratio. $\hat{\rho} = \mathrm{Corr}_{\mathrm{cs}}(\hat{\beta}_{i,p}, \hat{\alpha}_{i,p})$ is the cross-sectional correlation between full-sample betas and alphas from time-series regressions of the 1,277 sorted portfolios on the factor. $\text{OFB} = \hat{\gamma}_1^{\text{pop}} - \bar{r}_p$. Sharpe ratios are annualized.

STreversal ($|\hat{\rho}| = 0.78$, SR $= 1.51$) and DelFINL ($|\hat{\rho}| = 0.02$, SR $= 1.24$) have nearly identical Sharpe ratios but starkly different directional terms. STreversal has $|\hat{\rho}| = 0.78$ and a Sharpe ratio of 1.51. Its alphas are so aligned with its betas that the full-sample cross-sectional slope is 0.01, despite an expected return of 3.17%/month: the OFB almost perfectly offsets the risk premium. DelFINL has $|\hat{\rho}| = 0.02$ and a Sharpe ratio of 1.24. With near-zero directional misalignment, its full-sample slope of 0.49 is close to its expected return of 0.49%. The Sharpe ratio does not distinguish these cases; ρ does.

4.4 The Directional Component in Practice

The OFB formula predicts how the cross-sectional slope varies with the composition of the cross-section through $\mathrm{Cov}_{\mathrm{cs}}(\beta_{i,p},\ \alpha_{i,p})$. Different cross-sections produce different distributions of (β_i, α_i) pairs, and the formula maps each distribution to a specific slope.

Figure 5 illustrates this for SmileSlope, the factor with the highest absolute Sharpe ratio in the sample ($|SR| = 1.98$, $E[r_p] = 1.24\%$ per month). Using decile portfolios formed on six different characteristics, the estimated slope ranges from -5.15 (asset growth deciles) to $+1.92$ (accruals deciles). Using BM deciles, the slope is 1.25, close to $E[r_p]$; using momentum deciles, it is -1.51, wrong-signed and more than two standard errors from zero.

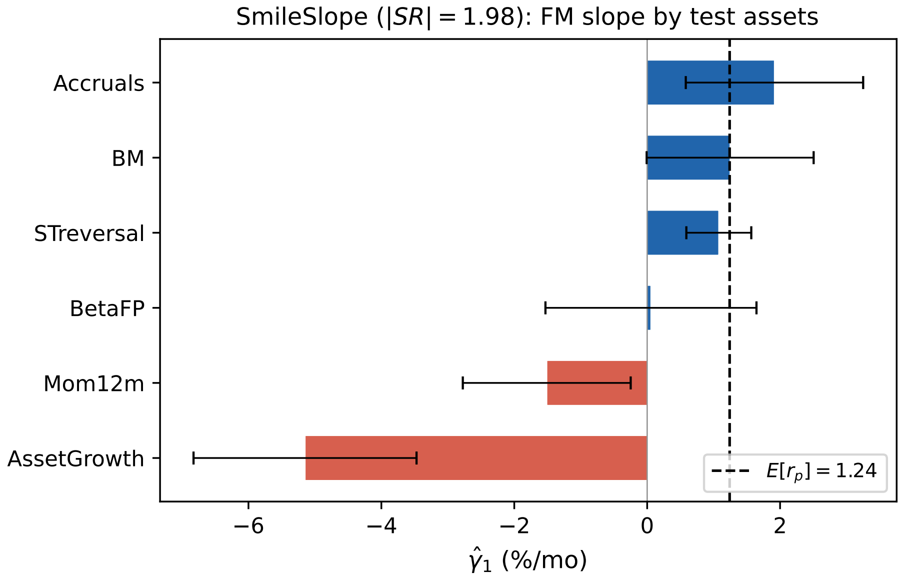

Figure 5: Fama-MacBeth slope for SmileSlope ($|SR| = 1.98$) using decile portfolios formed on six different characteristics. The dashed line marks $E[r_p] = 1.24\%$/mo. Error bars indicate one standard error. The same factor produces slopes ranging from -5.15 to $+1.92$ depending on the cross-section.

Table 5 extends this analysis to five factors spanning a wide range of Sharpe ratios. For each factor, the table reports the FM slope using all 1,277 portfolios as the baseline and the range across six decile-sort configurations. The range does not shrink with higher Sharpe ratios.

The OFB formula attributes this variation to changes in $\text{Cov}_{\text{cs}}(\beta_{i,p}, \alpha_{i,p})$. When the sort characteristic is correlated with an omitted SDF component, the cross-sectional covariance is large and the alignment term is large. When the sort is orthogonal to the omitted component, the covariance is small and γ_1 approaches $E[r_p]$.

Table 5: OFB variation across cross-sections

Factor	$\lvert SR \rvert$ (ann.)	$E[r_p]$ (%/mo)	Baseline γ_1 (All 1,277)	Min γ_1 (6 decile sorts)	Max γ_1
SmileSlope	1.98	1.24	0.11	−5.15	1.92
dCPVolSpread	1.64	1.14	0.18	−1.37	1.15
AssetGrowth	0.82	0.86	0.21	−1.19	1.00
BM	0.56	0.67	0.26	−1.71	3.30
GrAdExp	0.21	0.22	0.26	−5.20	4.39

Notes. Full-sample cross-sectional slopes (full-sample betas, single regression). Baseline uses all 1,277 sorted portfolios. Min and Max are over six decile-sort configurations (BM, Mom12m, AssetGrowth, Accruals, BetaFP, STreversal). Sharpe ratios are annualized.

Own-characteristic sorts amplify the alignment term. Sorting test assets on the characteristic underlying the factor maximizes $\text{Var}_{\text{cs}}(\beta_{i,p})$, the denominator of the OFB formula (19), because the sort is designed to spread betas. A larger denominator mechanically shrinks |OFB| for any given numerator. This reasoning, however, ignores the directional component. With $N \approx 10$ decile portfolios sorted on a single characteristic, the cross-section is effectively one-dimensional: both $\beta_{i,p}$ and $E[r_i]$ are smooth, monotonic functions of the decile number. The pricing error $\alpha_{i,p} = E[r_i] - \beta_{i,p}\, E[f_p]$ is then also smooth in the decile number, forcing $\rho_{\text{cs}}(\beta_{i,p}, \alpha_{i,p})$ toward ± 1. Formally, writing $E[r_i] = \gamma_1\, \beta_{i,p} + \gamma_0 + \epsilon_i$ across test assets, the pricing error becomes $\alpha_{i,p} = (\gamma_1 - E[f_p])\, \beta_{i,p} + \gamma_0 + \epsilon_i$. When the cross-sectional R^2 of $E[r_i]$ on $\beta_{i,p}$ is high, as it is for own-sort deciles, the residual ϵ_i is small and the alpha is dominated by the $(\gamma_1 - E[f_p])\, \beta_{i,p}$ term, pushing $|\rho|$ toward 1.

The mechanism also explains why ρ flips sign. With all 1,277 portfolios, the substitution structure documented in Section 2.11 produces $\rho < 0$ for 87% of factors: high-β_p assets tend to have low exposure to other priced factors. With own-sort deciles, this multi-dimensional substitution structure collapses. The 10 portfolios vary along a single characteristic, and the omitted factors' betas barely move across deciles. Instead, $\text{sign}(\rho) = \text{sign}(\gamma_1 - E[f_p])$, which is positive whenever the decile return spread exceeds the factor premium, the typical case because extreme deciles capture more of the characteristic's variation than the long-short portfolio.

Table 6 confirms this prediction across 185 factors with at least three own-sort portfolios. The median $|\text{OFB}/E[r_p]|$ rises from 0.87 with all 1,277 portfolios to 1.46 with own-sort deciles. The median $|\rho|$ jumps from 0.30 to 0.80, and the fraction of factors with $\rho < 0$ drops from 87% to 2.7%. Own-sort test assets produce smaller $|\text{OFB}|$ only 34.6% of the time: for the majority of factors, the directional alignment overwhelms the denominator advantage.

Table 6: Own-sort versus all-sort test assets

	Own-sort	All-sort
Test assets per factor	3–10	1277
Med. $|\text{OFB}/\overline{r}_p|$	1.46	0.87
Med. $|\rho_{\text{cs}}(\beta, \alpha)|$	0.80	0.30
Wrong sign	31.9%	30.8%
$\rho < 0$	2.7%	87.0%
Own-sort shrinks $|\text{OFB}|$	34.6%	

Notes. Comparison across 185 factors that have at least three own-characteristic sorted portfolios in Chen and Zimmermann (2022). "Own-sort" uses only the decile or quintile portfolios formed on the factor's underlying characteristic as test assets. "All-sort" uses all 1,277 sorted portfolios. Both columns use full-sample population betas and mean returns. $\rho_{\text{cs}}(\beta, \alpha)$ is the cross-sectional correlation between betas and pricing errors. Own-sort test assets increase $\text{Var}_{\text{cs}}(\beta)$ but collapse the cross-section to one dimension, pushing $|\rho|$ toward unity and amplifying the OFB for most factors.

The relevant distinction is not the number of test assets N but the effective dimensionality of the cross-section. Ten portfolios from ten different sorts would produce a multi-dimensional cross-section with moderate $|\rho|$; ten portfolios from one sort produce a one-dimensional cross-section with $|\rho|$ near unity. Maximizing beta spread is counterproductive when it comes packaged with maximum alpha-beta alignment.

4.5 Multi-Factor Models

The single-factor results extend to the multi-factor setting of Section 2.13. By Proposition 2, the full-sample cross-sectional slope for each factor in a multi-factor model decomposes into the factor's expected return plus an OFB that depends on the Frisch-Waugh-Lovell residual betas, not the raw betas. When factors are correlated, the FWL projection can amplify the

OFB through the variance ratio $\mathrm{Var}_{\mathrm{cs}}(\tilde{\beta}_k)/\,\mathrm{Var}_{\mathrm{cs}}(\tilde{\beta}_k^{\perp})$.

Table 7 reports results for three Fama-French factor models estimated against all 1,277 sorted portfolios: the three-factor model (FF3; Fama and French 1993), the five-factor model (FF5; Fama and French 2015), and the six-factor model (FF6, adding the UMD momentum factor). These are the most widely used factor models in the literature and the ones for which researchers routinely claim that high Sharpe ratios imply good cross-sectional pricing.

Table 7: Multi-factor Fama-MacBeth regressions

						FWL decomposition		
Model	Factor	γ_k^{pop}	$E[f_k]$	OFB	$\frac{\mathrm{OFB}}{E[f_k]}$	ρ_k	$\frac{\sigma(\alpha)}{\sigma(\tilde{\beta}_k^{\perp})}$	FWL
FF3	Mkt-RF	-0.622	0.595	-1.217	-2.04	-0.480	2.53	1.1
FF3	SMB	0.079	0.191	-0.111	-0.58	-0.119	0.93	1.1
FF3	HML	0.299	0.271	0.028	0.10	0.023	1.20	1.0
FF5	Mkt-RF	-0.257	0.582	-0.839	-1.44	-0.296	2.83	1.3
FF5	SMB	0.131	0.186	-0.054	-0.29	-0.048	1.13	1.3
FF5	HML	0.164	0.265	-0.101	-0.38	-0.063	1.60	1.7
FF5	RMW	0.116	0.269	-0.153	-0.57	-0.098	1.57	2.1
FF5	CMA	0.335	0.169	0.166	0.98	0.089	1.86	1.4
FF6	Mkt-RF	-0.193	0.582	-0.775	-1.33	-0.288	2.69	1.1
FF6	SMB	0.209	0.186	0.023	0.12	0.021	1.10	1.4
FF6	HML	0.422	0.265	0.157	0.59	0.079	1.99	2.1
FF6	RMW	-0.009	0.269	-0.279	-1.04	-0.174	1.60	2.4
FF6	CMA	0.247	0.169	0.078	0.46	0.043	1.81	1.2
FF6	UMD	0.544	0.583	-0.039	-0.07	-0.019	2.09	1.4

Notes. Each row reports results for one factor within a multi-factor Fama-MacBeth regression against all 1,277 sorted portfolios. γ_k^{pop} is the population cross-sectional slope (full-sample betas, single OLS cross-sectional regression). OFB $= \gamma_k^{\mathrm{pop}} - E[f_k]$; OFB$/E[f_k]$ expresses the bias as a fraction of the factor mean. The FWL decomposition follows Proposition 2: $\mathrm{OFB}_k = \rho_k \cdot \sigma_{\mathrm{cs}}(\alpha)/\sigma_{\mathrm{cs}}(\tilde{\beta}_k^{\perp})$, where $\rho_k = \rho_{\mathrm{cs}}(\tilde{\beta}_k^{\perp}, \alpha)$ is the cross-sectional correlation between the FWL-residualized beta and the pricing error, and $\sigma(\alpha)/\sigma(\tilde{\beta}_k^{\perp})$ is the variance amplification from projecting out the other factors. FWL $= \mathrm{Var}_{\mathrm{cs}}(\tilde{\beta}_k)/\,\mathrm{Var}_{\mathrm{cs}}(\tilde{\beta}_k^{\perp})$ measures the multicollinearity inflation ratio.

The OFB$/E[f_k]$ column expresses the alignment term as a fraction of the factor's expected return. The market factor's alignment term exceeds its expected return in magnitude and flips the sign of γ_1 in all three models (-2.04 in FF3, -1.33 in FF6). RMW worsens from -0.57 in FF5 to -1.04 in FF6, crossing from attenuated to wrong-signed. UMD's ratio is

-0.07: the alignment term is 7% of its expected return. The second component is not a uniform nuisance; it ranges from negligible to dominant within the same model, and the Sharpe ratio does not predict which factors are affected.

The market factor has a wrong-sign full-sample slope in all three models: $\hat{\gamma}^{\text{fs}} = -0.622$ in FF3, -0.257 in FF5, and -0.193 in FF6, against $E[\text{Mkt-RF}] = 0.58$–0.60%/month. The cross-sectional regression assigns a negative price of risk to market beta across all specifications: the empirical counterpart of the flat security market line and the betting-against-beta pattern. Adding factors reduces the market's OFB (from -1.22 to -0.78) but does not eliminate it.

Comparing FF5 to FF6 illustrates how adding a single factor reshuffles individual OFBs. The RMW slope drops from 0.116 in FF5 to -0.009 in FF6, a wrong-sign estimate, even though its expected return (0.269%/month) is unchanged. The HML slope moves in the opposite direction, from 0.164 to 0.422, overshooting its expected return. Adding UMD improves the model's aggregate fit (R^2 rises from 0.136 to 0.188) but worsens the OFB for RMW (from -0.15 to -0.28) and flips the sign of the OFB for HML (from -0.10 to $+0.16$).

The FWL decomposition in the right columns of Table 7 traces these shifts to two quantities: ρ_k, the cross-sectional correlation between factor k's FWL-residualized betas and the model's pricing errors, and $\sigma(\alpha)/\sigma(\tilde{\beta}_k^{\perp})$, the ratio of pricing-error dispersion to residualized-beta dispersion. The market factor has the largest OFB because both components work against it: a strongly negative ρ (-0.48 in FF3, -0.29 in FF6) combines with a high variance ratio (2.5–2.8). By contrast, UMD has a near-zero ρ (-0.02), so despite a high variance ratio (2.1), its OFB is small (-0.04). The FWL column reports the multicollinearity inflation ratio $\operatorname{Var}_{\text{cs}}(\tilde{\beta}_k)/\operatorname{Var}_{\text{cs}}(\tilde{\beta}_k^{\perp})$: values range from 1.0 to 2.4, indicating that factor collinearity amplifies the alignment term modestly. The instability of RMW between FF5 and FF6 traces to a rise in its FWL ratio from 2.1 to 2.4 combined with a deeper ρ (from -0.10 to -0.17): adding UMD shrinks RMW's residualized-beta variance and strengthens the correlation between the residualized betas and the pricing errors.

The standard Fama-MacBeth estimates $\hat{\gamma}_k$ track the full-sample slopes, confirming that the alignment term, not estimation noise, drives the gap between γ_k and $E[f_k]$. The multi-factor results reinforce the single-factor finding: even in canonical models with well-studied factors and full-sample betas (eliminating estimation error by construction), the FM slope does not coincide with the factor's expected return.

4.6 Calibrating the Cauchy-Schwarz Bound

Proposition 1 provides a Cauchy-Schwarz bound on the alignment term that tightens as the factor's Sharpe ratio approaches the tangency portfolio's Sharpe ratio. The tangency Sharpe ratio is not directly observable, but it is bounded below by the maximum Sharpe ratio achievable from any set of traded factors, providing a range of plausible values. The question is whether this bound is tight enough, for realistic parameter values, to rule out wrong-sign slopes. If so, the Sharpe ratio would constrain γ_1 even under misspecification; if not, the bound is loose and the Sharpe ratio is practically uninformative about γ_1.

The bound states $|\gamma_1 - E[r_p]| \leq C \cdot \sqrt{\mathrm{SR}^2(\text{tangency}) - \mathrm{SR}^2(p)}$, where $C = \sqrt{\tilde{\beta}'\Sigma_e\tilde{\beta}}/(\tilde{\beta}'\tilde{\beta})$ is a factor- and test-asset-specific prefactor. A wrong-sign slope is possible whenever the bound exceeds $|E[r_p]|$. For each of the 212 factors, I compute the prefactor C using the full-sample betas and residual covariance matrix from the 1,277 test assets, then evaluate the bound for a range of SR(tangency) values.

Empirical estimates of SR(tangency) span a wide range. Standard multi-factor models produce in-sample annualized maximum Sharpe ratios of 1.0–2.4: roughly 0.6 for the Fama-French three-factor model, 1.2 for the five-factor model, and 2.0–2.4 for six-factor models and the q^5 model (Barillas et al. 2020; Fama and French 2018). Easterwood and Paye (2025) show that much of this reflects data snooping: using validation samples that did not influence factor construction, the maximum Sharpe ratio of the FF6 model drops from approximately 2.4 to 1.2. Machine-learning approaches report out-of-sample Sharpe ratios as high as 2.5 (Kelly, Pruitt, and Su 2019). I evaluate the bound over SR(tangency) $\in$

Table 8: Calibrating the Cauchy-Schwarz bound

SR(tangency)	Frac. wrong sign permitted	Median max $\lvert\text{OFB}/E[r_p]\rvert$	P75 max $\lvert\text{OFB}/E[r_p]\rvert$	P90 max $\lvert\text{OFB}/E[r_p]\rvert$
Panel A: OLS bound (Proposition 1)				
1.0	0.83	2.3	4.6	9.9
1.5	0.95	3.8	7.5	15.0
2.0	0.99	5.3	10.2	20.0
2.5	1.00	6.7	12.9	25.0
3.0	1.00	8.1	15.6	30.0
4.0	1.00	10.9	20.9	40.1
5.0	1.00	13.7	26.2	50.1
Panel B: GLS bound (Remark 4)				
1.0	0.16	0.4	0.6	1.3
1.5	0.25	0.6	1.0	2.0
2.0	0.38	0.8	1.3	2.7
2.5	0.54	1.1	1.6	3.4
3.0	0.64	1.3	2.0	4.1
4.0	0.82	1.7	2.7	5.4
5.0	0.92	2.2	3.3	6.8

Notes. For each of 212 factors, the Cauchy-Schwarz bound gives the maximum $\lvert\gamma_1 - E[r_p]\rvert$ consistent with the factor's Sharpe ratio and a given SR(tangency). Panel A uses the OLS prefactor $\sqrt{\tilde{\beta}'\Sigma_e\tilde{\beta}}\,/\,(\tilde{\beta}'\tilde{\beta})$ from Proposition 1. Panel B uses the GLS prefactor $1/\sqrt{\tilde{\beta}'\Sigma_e^{-1}\tilde{\beta}}$ from Remark 4, with a diagonal $\hat{\Sigma}_e$ (residual variances from the time-series regressions). "Frac. wrong sign permitted" is the share of factors for which the bound exceeds $\lvert E[r_p]\rvert$. All Sharpe ratios are annualized. Empirical estimates place SR(tangency) in the range 1.0–2.5 (Easterwood and Paye 2025; Barillas et al. 2020).

$\{1.0, 1.5, 2.0, 2.5, 3.0, 4.0, 5.0\}$ to span conservative through aggressive estimates.

Table 8 reports the results. Panel A uses the OLS prefactor from Proposition 1. Even at the most conservative value of SR(tangency) = 1.0, the OLS bound permits wrong-sign slopes for 83% of the 212 factors: the bound on $\lvert\gamma_1 - E[r_p]\rvert$ exceeds $\lvert E[r_p]\rvert$ for 83% of them. The median maximum $\lvert\text{OFB}/E[r_p]\rvert$ permitted by the bound is 2.3 at SR(tangency) = 1.0 and 5.3 at SR(tangency) = 2.0.

Panel B uses the GLS prefactor from Remark 4, $1/\sqrt{\tilde{\beta}'\Sigma_e^{-1}\tilde{\beta}}$, with a diagonal $\hat{\Sigma}_e$ estimated from the time-series regressions. The GLS bound is substantially tighter: at SR(tangency) = 1.0, the GLS bound permits wrong-sign slopes for only 16% of factors, compared to 83% under OLS. The median maximum $\lvert\text{OFB}/E[r_p]\rvert$ drops from 2.3 to 0.37.

The improvement comes from the metric alignment between the GLS estimator and the GRS quadratic form (Remark 4): the OLS bound suffers from a mismatch between the Euclidean norm in the estimator and the Σ_e^{-1} norm in the GRS identity, which the GLS bound eliminates. At SR(tangency) = 2.0, the GLS bound still permits wrong-sign slopes for 38% of factors, and the observed wrong-sign rate under population GLS is 23% (Table 3, Panel B), within the bound. The GLS bound disciplines the alignment term more effectively than the OLS bound but does not eliminate the problem: a nontrivial fraction of factors can still have wrong-sign GLS slopes at realistic tangency Sharpe ratios.

The comparison between the observed GLS wrong-sign rate and the bound is informative in both directions. The observed population GLS wrong-sign rate is 23% (Table 3, Panel B), which exceeds the 16% permitted by the GLS bound at SR(tangency) = 1.0. The bound is a necessary condition: if the observed wrong-sign rate exceeds it, the assumed tangency Sharpe ratio is too low. The GLS bound therefore implies SR(tangency) > 1.0, consistent with standard multi-factor model estimates (Barillas et al. 2020). At SR(tangency) = 2.0, the GLS bound permits 38%, comfortably above the observed 23%.

Inverting the GLS bound. The GLS bound can be inverted to ask: for each factor p, what is the minimum tangency Sharpe ratio that permits a wrong-sign slope? Solving the GLS wrong-sign condition $|E[r_p]| \leq (1/\sqrt{\tilde{\beta}'\Sigma_e^{-1}\tilde{\beta}}) \cdot \sqrt{\mathrm{SR}^2(\text{tangency}) - \mathrm{SR}^2(p)}$ for SR(tangency):

$$\mathrm{SR}_{\min}(p) = \sqrt{\mathrm{SR}^2(p) + E[r_p]^2 \cdot \tilde{\beta}'\Sigma_e^{-1}\tilde{\beta}}. \tag{50}$$

Table 9 reports the cross-factor distribution of $\mathrm{SR}_{\min}(p)$. The median is 2.35: for the typical factor, the tangency portfolio must have an annualized Sharpe ratio above 2.35 before the GLS bound even permits a wrong-sign slope. At the conservative end, 84% of factors require SR(tangency) > 1.0. These numbers use the diagonal $\hat{\Sigma}_e$ approximation and are therefore lower bounds on the full-covariance $\mathrm{SR}_{\min}$; the full Σ_e^{-1} would increase $\tilde{\beta}'\Sigma_e^{-1}\tilde{\beta}$ and tighten the requirement. The inversion complements the GRS-based estimate $\mathrm{SR}^2(\text{tangency})$ =

Table 9: Implied lower bound on the tangency Sharpe ratio

Panel A: Distribution of $\mathrm{SR}_{\min}(p)$ *across 212 factors*	
Min	0.04
P25	1.55
Median	2.35
P75	3.53
P90	4.80
Max	12.64

Panel B: Fraction of factors requiring $\mathrm{SR}(\tan) >$ *threshold*	
1.0	84.4%
1.5	75.5%
2.0	61.8%
2.5	45.8%
3.0	35.8%

Notes. $\mathrm{SR}_{\min}(p)$ is the minimum annualized tangency Sharpe ratio for which the GLS Cauchy-Schwarz bound (Remark 4) permits a wrong-sign slope for factor p: $\mathrm{SR}_{\min}(p) = \sqrt{\mathrm{SR}^2(p) + E[r_p]^2 \cdot \tilde{\beta}'\Sigma_e^{-1}\tilde{\beta}}$, with diagonal $\hat{\Sigma}_e$. Panel A reports the cross-factor distribution. Panel B reports the fraction of factors whose $\mathrm{SR}_{\min}(p)$ exceeds each threshold. A median of 2.35 means that for the typical factor, the tangency portfolio must have an annualized Sharpe ratio above 2.35 before the bound even permits a wrong-sign slope.

$\mathrm{SR}^2(p) + \alpha'\Sigma_e^{-1}\alpha$, which gives a direct (and tighter) lower bound but requires estimating the full quadratic form $\alpha'\Sigma_e^{-1}\alpha$, infeasible with a singular Σ_e.

The tighter GLS bound reflects a less loose inequality, not necessarily a smaller alignment term. The OLS bound is loose because the OLS estimator weights assets by $\tilde{\beta}^2$ while the GRS identity weights them by Σ_e^{-1}; this metric mismatch means the Cauchy-Schwarz inequality is far from attained. The GLS estimator and the GRS identity share the Σ_e^{-1} metric, so the bound is closer to binding. But the actual GLS OFB, $\tilde{\beta}'\Sigma_e^{-1}\alpha/(\tilde{\beta}'\Sigma_e^{-1}\tilde{\beta})$, is a different inner product from the OLS OFB, $\tilde{\beta}'\alpha/(\tilde{\beta}'\tilde{\beta})$: GLS reweights assets by the inverse residual variance, changing which assets drive the slope. This reweighting can increase or decrease the bias depending on whether high-precision assets happen to concentrate or disperse the correlation between β and α. The observed population GLS wrong-sign rate (23%, Table 3) is slightly higher than the OLS rate (21%), confirming that the tighter bound does not imply a smaller actual alignment term. The GLS bound constrains the worst case more tightly, but the actual GLS slope lives in a different part of the feasible set.

Figure 6 shows the empirical counterpart. The wrong-sign rate drops from 44% in the lowest $|SR|$ quintile to 2% in the highest. By Remark 6, both the Sharpe ratio and the wrong-sign condition are scale-invariant, so this pattern is not a threshold artifact from larger $|E[r_p]|$: it reflects that high-Sharpe-ratio factors have covariance structures better aligned with the cross-section of expected returns (D^* tends to agree in sign with $E[r_p]$).

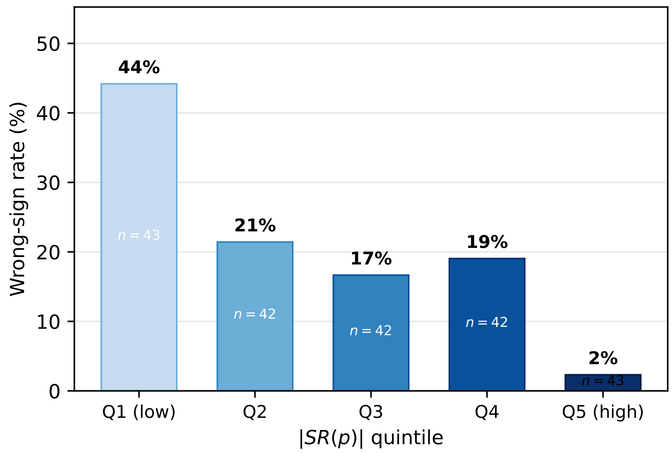

Figure 6: Wrong-sign rate by $|SR(p)|$ quintile for 212 long-short factors. A wrong sign means the full-sample cross-sectional slope γ_1^{pop} and the factor's expected return $E[r_p]$ have opposite signs. The rate drops from 44% (Q1) to 2% (Q5): the Sharpe ratio predicts the sign of γ_1 but not its magnitude.

5 Implications for Practice

The Fama-MacBeth slope has two components: the factor's expected return and a cross-sectional alignment term. The Sharpe ratio constrains the first but not the second, wrong-sign slopes are common, and no estimator eliminates the alignment term. No estimator *can* eliminate it: the alignment term is a population feature of the misspecified model, not an estimation artifact. The pricing equation (4) shows that $E[r_p]$ is the coefficient on $\beta_{i,p}$ only conditional on controlling for $\text{Cov}(\varepsilon, r_i)/E[m]$, the contribution of every SDF component not spanned by r_p. The Fama-MacBeth regression cannot include this term because ε is unobservable, and no weighting scheme removes it (Corollary 3). Even $E[r_p]$ itself, the factor's expected return, cannot be interpreted as "the price of bearing factor-p beta risk" without the maintained hypothesis that the factor belongs to the true SDF. If the factor is outside the SDF but correlated with priced factors, $E[r_p]$ reflects borrowed risk premia from those factors, and "the price of factor-p beta risk" is not a primitive concept in the true pricing kernel. Under misspecification, the two components of γ_1 differ, and neither alone constitutes a model-free price of risk; only correct specification makes the cross-sectional slope equal to the factor's return and gives it an unambiguous economic interpretation. Testing whether the factor belongs to the true SDF is itself either trivial or impossible: the univariate condition ($b \neq 0$, equivalently $E[r_p] \neq 0$) is satisfied by every factor with a nonzero mean, while the multivariate condition ($b_p \neq 0$ in the true SDF) requires knowledge of the complete pricing kernel.

The Fama-MacBeth regression as a portfolio device. Section 2.5 establishes that the cross-sectional slope γ_{1t} is the return on a zero-cost, unit-beta mimicking portfolio with weights given by (6). The time-series average $\bar{\gamma}_1$ is a consistent estimator of that portfolio's expected return. No SDF assumption is needed for this statement: the weights are mechanical functions of estimated betas, the portfolio return is an observable cash flow, and the law of large numbers delivers consistency. As a portfolio formation device, the Fama-MacBeth regression is unimpeachable.

The problem is the next step: interpreting $\bar{\gamma}_1$ as "the price of risk" and expecting it to equal $E[r_p]$. That requires the alignment term to vanish, which by (8) requires OFB $= 0$. The alignment term is a population quantity, the cross-sectional regression coefficient of alphas on betas (9), and it does not shrink with T. Increasing the sample sharpens the estimate of a quantity whose two components generically differ: the researcher converges to γ_1, a composite of $E[r_p]$ and the alignment term, not to $E[r_p]$ alone. Kan and Zhang (1999) show that inference fails for irrelevant factors: the t-statistic diverges even when the true beta is zero. The decomposition shows that the point estimate is ambiguous for relevant factors: γ_1 reflects the factor's return and the cross-sectional alignment, and no amount of data separates the two. The Fama-MacBeth procedure constructs portfolios; the slope is the return on a specific portfolio, not a measure of what beta is worth in the cross-section.

Separate model comparison from pricing inference. The maximum squared Sharpe ratio is the correct metric for comparing factor models (Barillas and Shanken 2017; Barillas et al. 2020): the model whose factors achieve a higher SR^2 produces smaller aggregate alphas. But a model that wins this comparison need not produce Fama-MacBeth slopes equal to the factors' expected returns. The two questions, "which model prices assets better in aggregate?" and "does this factor's beta command a cross-sectional premium equal to $E[f_k]$?", have different answers in general. The first depends on $\|\alpha\|$; the second depends on $\mathrm{Cov}_{\mathrm{cs}}(\beta, \alpha)$. Researchers should state which question they are addressing and should not cite a high Sharpe ratio as evidence for the second.

The three-pass estimator and the limits of PCA. Giglio and Xiu (2021) propose a three-pass procedure that uses principal components of test-asset returns as instruments for the unobserved latent factors. Their cross-sectional regression of expected returns on PCA betas delivers slopes equal to the factors' expected returns under two key assumptions: (i) the latent factors that drive expected returns are the same factors that drive return covariation (their Assumption A1), and (ii) these factors are "strong," meaning their contribution to return variance does not vanish as the number of test assets grows (Assumption A6, pervasiveness). When both conditions hold, the PCA factors span the true pricing kernel and the three-pass estimator is consistent for the risk prices, a genuine advance over Fama-MacBeth. The maintained hypothesis, that dominant PCA factors span the SDF, may be substantially more plausible than the hypothesis that any particular set of researcher-chosen factors does.

The decomposition clarifies the boundary of this advance. If the omitted SDF component ε in (1) is driven by a weak factor, one with a nonzero price of risk but low variance contribution, PCA will not extract it: the factor's eigenvalue is too small to distinguish from noise. The three-pass estimator then omits the same factor that the original model omits, and γ_1 retains the same alignment term. Giglio, Xiu, and Zhang (2021) acknowledge this limitation and propose penalized estimators for the weak-factor case, but these require additional tuning parameters and identification assumptions. The three-pass procedure makes the two components of γ_1 coincide when PCA spans the SDF; it does not when a priced factor is weak. The alignment term is a population feature of misspecification, and it persists whenever the included factors, whether chosen by the researcher or extracted by PCA, fail to span the true pricing kernel.

Report the directional diagnostic $\hat{\rho}$. The decomposition in Corollary 1 identifies $\rho_{\text{cs}}(\beta, \alpha)$ as the component of the OFB that the Sharpe ratio does not constrain. This quantity is computable from the same time-series regressions that produce betas and alphas. I recommend that researchers report $\hat{\rho}$ alongside the Fama-MacBeth slope and its standard error. A large $|\hat{\rho}|$ signals that the pricing errors are directionally aligned with betas, meaning the alignment term dominates γ_1. The wrong-sign condition (Corollary 2) provides a direct check: if $|\hat{\rho}| \cdot \text{SD}_{\text{cs}}(\hat{\alpha})/\text{SD}_{\text{cs}}(\hat{\beta}) > |E[r_p]|$, the estimated slope will have the wrong sign. Unlike the Sharpe ratio, $\hat{\rho}$ speaks directly to the cross-sectional regression and can flag problems that the Sharpe ratio misses.

Report sensitivity to test assets and model specification. The alignment term depends on the test assets through $\text{Cov}_{\text{cs}}(\beta_{i,p}, \alpha_{i,p})$, so γ_1 is not a property of the factor alone. I recommend that researchers report the FM slope across at least two or three standard test-asset configurations (e.g., size/value 25, industry 48, broad characteristic sorts). If γ_1 is stable across configurations, the alignment term is likely small or at least consistently signed. If γ_1 varies substantially, the interaction between the factor and the specific cross-section, not the factor's own return, drives the result. The check requires only standard data and would have flagged many of the failures documented in Section 4.

In multi-factor models, the FWL structure (Proposition 2) makes the OFB for each factor sensitive to which other factors are included. Adding a factor changes the FWL residual betas for every existing factor, reshuffling the individual alignment terms even when the model's aggregate fit improves. The RMW slope flips sign between FF5 and FF6 (Table 7), illustrating that γ_k can depend as much on the other factors in the model as on the test assets. Reporting slopes across model specifications, not just test-asset configurations, would make this instability visible.

Time-varying betas. The OFB formula applies to each cross-section separately, so the analysis extends to conditional models with time-varying betas and risk premia. The empirical results in Table 3 already use time-varying betas: the rolling 60-month and daily (252-day, 1260-day) specifications allow betas to change over time, and the Fama-MacBeth regression runs period by period with those betas. The patterns are similar across all four beta methods: wrong-sign rates range from 22% to 31%, and the correlation between $|SR|$ and the FM slope never exceeds 0.19. Allowing betas to vary adds estimation noise (the attenuation factor λ falls from 0.88 for daily 1260-day betas to 0.56 for rolling 60-month betas) but does not resolve the OFB, because the OFB is a property of the population cross-section, not of the beta estimator. The Fama-MacBeth procedure estimates the time-series average of the period-by-period slope, $E[\gamma_{1t}] = E[E_t[r_p] + \text{OFB}_t]$, which accommodates time-varying risk premia by averaging over them. If betas and risk premia covary over time, the unconditional FM slope reflects $E[\beta_{it}\lambda_t]$ rather than $E[\beta_{it}] \cdot E[\lambda_t]$, a standard distinction between conditional and unconditional models. This distinction is orthogonal to the OFB: even a correctly conditional model produces an OFB unless the factors span the SDF in every period.

Constructive use of the decomposition. The decomposition is not only a warning; it suggests when the FM slope is likely to be informative. The key is the effective dimensionality of the cross-section. A multi-dimensional set of test assets, sorted on many characteristics, allows the pricing errors to vary in directions orthogonal to beta, keeping

$|\rho_{\text{cs}}(\beta, \alpha)|$ moderate. A one-dimensional set, such as own-sort deciles, forces $|\rho|$ toward unity (Section 4, Table 6): the common practice of using own-sort portfolios to maximize beta spread is counterproductive because it packages maximum beta dispersion with maximum alpha-beta alignment. The broader the cross-section, the more room for the pricing errors to scatter in directions orthogonal to beta. When the FM slope is stable across test-asset configurations and $|\hat{\rho}|$ is small, the researcher can have reasonable confidence that γ_1 reflects the factor's return rather than the alignment term. The decomposition does not say that the FM regression is always uninformative; it says that whether γ_1 reflects the factor's return depends on a verifiable condition, $\text{Cov}_{\text{cs}}(\beta, \alpha) \approx 0$, rather than on the Sharpe ratio.

5.1 Out-of-sample implications

The OFB formula is an algebraic identity: it holds with $R^2 = 1$ for every factor when population betas are used. Identities are not testable. The decomposition nevertheless generates falsifiable predictions about patterns across factors, test assets, and model specifications that go beyond the identity itself.

Subsample stability. The OFB identity holds exactly in any sample. Whether the cross-sectional covariance structure it describes is a durable feature of the economy is a separate, empirical question. Table 10 splits the sample into three 20-year subperiods (1963–1983, 1984–2004, 2005–2024) and recomputes full-subsample betas, mean returns, and the OFB independently within each window.

Panel A shows that the OFB phenomenon persists in every era: wrong-sign rates range from 30% to 39%, higher than the full-sample 21%, and the median $|\text{OFB}/\bar{r}_p|$ exceeds 0.89 in all three subperiods. The fraction of factors with $\rho < 0$ ranges from 32% to 52%, well below the full-sample 87%. This drop reflects measurement error attenuation: within a 20-year window, the sampling noise in each asset's mean return ($\sigma_i/\sqrt{T} \approx 0.5\%$/month) is comparable to the cross-sectional dispersion of true pricing errors, attenuating the measured $|\rho|$ toward zero. The higher wrong-sign rates confirm that shorter estimation windows compound the problem rather than resolve it.

Panel B asks whether the same factors have large OFBs across eras. Adjacent subperiods show moderate, highly significant rank stability: Spearman correlations of 0.40 for the OFB and 0.38–0.41 for ρ, with 62–65% of factors retaining the same OFB sign. Distant subperiods (1963–1983 vs. 2005–2024) show no significant rank correlation ($r_s = 0.11$, $p = 0.14$), and the OFB sign is a coin flip (51%). The structural ρ equation (35) explains this decay: ρ depends on the risk-premium-weighted cross-sectional beta covariances with omitted factors, and both the risk premia $E[f_k]$ and the covariance structure $\text{Cov}_{\text{cs}}(\beta_p, \beta_k)$ evolve over 40-year horizons as the factor structure of the economy changes. The OFB phenomenon appears in every subperiod; which specific factors bear the largest bias does not persist across distant eras.

Table 10: Subsample stability of OFB

Panel A: Per-subperiod statistics

Period	K	Wrong sign	Med. $\lvert\text{OFB}/\bar{r}_p\rvert$	$\rho < 0$	$\lvert\rho\rvert$
1963-1983	173	39.3%	1.44	31.8%	0.324
1984-2004	211	37.0%	0.89	47.9%	0.232
2005-2024	209	30.1%	1.17	52.2%	0.199

Panel B: Cross-period rank correlations (Spearman)

Period pair	K	r_s(OFB)	p-value	$r_s(\rho)$	Same sign
1963-1983 vs 1984-2004	173	0.396***	0.0000	0.375	62.4%
1963-1983 vs 2005-2024	173	0.113	0.1398	0.096	50.9%
1984-2004 vs 2005-2024	208	0.392***	0.0000	0.406	65.4%

Notes. Panel A reports OFB statistics for each 20-year subperiod. Within each subperiod, population betas and mean returns are estimated using only that subperiod's data. K is the number of factors with sufficient data (at least 60 months). "Wrong sign" is the fraction of factors where γ_1^{pop} and $\bar{r}_p$ have opposite signs. Panel B reports Spearman rank correlations of the OFB and $\rho_{\text{cs}}(\beta, \alpha)$ across subperiods for factors present in both. "Same sign" is the fraction of factors whose OFB has the same sign in both subperiods. ***, **, * denote significance at the 1%, 5%, 10% levels.

An observable proxy for ρ: promise and pitfalls. Equation (35) expresses ρ in terms of the risk-premium-weighted cross-sectional beta covariances with omitted factors. While the true omitted factors are unknown, the observed factor zoo serves as a proxy. Define

$$\widehat{\mathrm{OFB}}_p = \sum_{k \neq p} E[f_k] \cdot \frac{\mathrm{Cov}_{\mathrm{cs}}(\beta_{i,p},\, \beta_{i,k})}{\mathrm{Var}_{\mathrm{cs}}(\beta_{i,p})},$$

where all betas are univariate. The formula is exact when factors are independent; with correlated factors it is an approximation because univariate betas confound the true exposure to factor k with exposures to every correlated factor. A subtler problem is near-tautology. If the factor zoo approximately spans the cross-section of expected returns, $E[r_i] \approx \sum_k c_k\, \beta_{i,k}^{\mathrm{univ}}$, then both the proxy and the actual OFB reduce to weighted sums of $\mathrm{Cov}_{\mathrm{cs}}(\beta_{i,p}, \beta_{i,k}) / \mathrm{Var}_{\mathrm{cs}}(\beta_{i,p})$ with different weights ($E[f_k]$ versus c_k). The resulting in-sample correlation may be mechanical rather than informative. A non-trivial test would require a time-split design (proxy from the first half, OFB from the second) or a horse-race against simpler alternatives such as the first few principal components of the factor zoo. Whether the structural formula adds predictive content beyond what a low-dimensional factor model already provides remains an open question.

Predicting the direction of cross-model OFB changes. Adding factor f_{K+1} to a K-factor model shifts the pricing errors from α_i^{old} to approximately $\alpha_i^{\mathrm{old}} - \beta_{i,K+1}\, \gamma_{K+1}$, removing the component of mispricing aligned with $\beta_{i,K+1}$. The structural equation (35) predicts the *sign* of the resulting OFB change: when $\mathrm{Cov}_{\mathrm{cs}}(\tilde{\beta}_{i,k}^{\perp},\, \beta_{i,K+1})$ has the same sign as the current OFB_k and the added factor carries a positive premium, the addition should move γ_k closer to $E[f_k]$; when the signs oppose, the OFB should worsen. The RMW sign flip between FF5 and FF6 (Table 7) is one instance.

Mechanical tests and the identity boundary. Because the OFB formula is an algebraic identity, several seemingly natural empirical tests turn out to be mechanical restatements rather than falsifiable predictions. Recognizing this boundary matters for directing future empirical work. Three examples illustrate the pattern. First, the intercept-slope seesaw

($\gamma_0 = \bar{\mu} - \gamma_1\bar{\beta}$) is a property of OLS, not of the OFB: the intercept absorbs whatever the slope does not, in every cross-sectional regression, misspecified or not. Second, comparing OLS and GLS sign disagreements restates when two weighted sums of the same alpha vector have opposite signs (Corollary 3), adding no content beyond the universality result. Third, drawing random subsets of test assets of increasing size and checking whether the FM slope converges away from $E[r_p]$ is not a law-of-large-numbers experiment but a combinatorial resampling exercise on a finite set; every stable cross-sectional statistic converges under such resampling, so the test has no power. Similarly, the *magnitude* of ΔOFB_k when adding a factor is largely mechanical: the cross-sectional R^2 of $\beta_{i,k}$ on $\beta_{i,K+1}$ determines how much the FWL denominator shrinks, a generic property of multivariate regression. The general lesson is that any "prediction" derived purely from the algebra of the identity, without additional economic structure, will hold by construction and therefore cannot be refuted by data.

6 Conclusion

The stochastic discount factor exists as a consequence of no-arbitrage. No particular set of traded factors composes it; it is a function of the entire state of the economy. Any set of portfolios can serve as factors, and the SDF can always be projected onto them, but the projection is not the SDF. The residual ε from this projection is the component of the SDF that lives outside the span of the chosen factors, and it is present whenever the model is misspecified. The Sharpe ratio of the factor portfolio measures the share of SDF variance the projection captures. The entire logic of the Sharpe-ratio-based paradigm in empirical asset pricing, from model comparison (Barillas and Shanken 2017; Barillas et al. 2020) to factor design (Kozak, Nagel, and Santosh 2018, 2020), rests on the implicit assumption that a factor that earns a high return (high Sharpe ratio) must also price the cross-section well (Fama-MacBeth slopes close to the factor's return). The assumption is wrong: earning a return and pricing the cross-section are different properties, connected only under spanning.

Writing $E[r_i] = \beta_i'\lambda$ is not a statistical model for estimation; it is correct specification of the SDF. It asserts that the projection has no residual ($\varepsilon = 0$), that the factors span the SDF completely, and that betas on those factors are sufficient statistics for expected returns. Under this maintained hypothesis, the Sharpe ratio and the Fama-MacBeth slope carry the same information. But empirical work proceeds without knowing whether the specification is correct, and once any residual exists, the connection breaks. The standard derivation begins from $E[r_i] = \beta_i'\lambda$ and assumes the two components of γ_1 are identical before anyone can examine whether they are; starting from the SDF projection reveals that γ_1 depends on both the factor's return and a cross-sectional alignment term that the factor's own properties do not determine. The Sharpe ratio constrains the total variance of the SDF residual through the Hansen-Jagannathan bound, but not how that residual distributes across the cross-section of betas, so the alignment term is loosely bounded in theory and practically unrestricted for realistic parameter values.

The two components of γ_1 coincide for all test assets if and only if the model is correctly specified, meaning the factors span the SDF (Corollary 3). For a single fixed cross-section, the alignment term also vanishes when pricing errors happen to be orthogonal to betas (Corollary 1(b)), but this condition is knife-edge and unverifiable without knowing the true model. No estimator eliminates the alignment term: GLS inherits the same structure as OLS, and the three-pass procedure of Giglio and Xiu (2021) is consistent under the assumption that the factors driving return covariation are exactly the factors that determine expected returns and are pervasive enough for PCA to recover (their Assumptions A1 and A6). No diagnostic based on the Sharpe ratio, or on any other function of the included factors alone, can determine the alignment term or provide a practically binding bound on it. The factor's expected return is a real cash flow, but whether its beta carries a cross-sectional price equal to that return, meaning that assets with higher beta earn proportionally more, depends on whether the factor has marginal pricing power in the true SDF, a condition that cannot be verified without observing the complete pricing kernel.

Across 212 long-short factors from the Chen and Zimmermann (2022) database, the correlation between Sharpe ratios and Fama-MacBeth slopes never exceeds 0.19, wrong-sign slopes are common even when errors-in-variables attenuation is minimal, and the highest Sharpe ratio factor produces slopes ranging from -5.15 to $+1.92$ depending solely on the choice of test assets. Multi-factor models exhibit the same disconnect: the Frisch-Waugh-Lovell structure amplifies the alignment term when factor betas are correlated across the cross-section, so that adding factors can increase the gap between what individual factors earn and what their betas price, even as the model's aggregate Sharpe ratio improves.

The Sharpe ratio remains the right metric for the question it was designed to answer: which set of factors captures a larger share of the SDF? But the field has treated a factor's return and its cross-sectional price as the same object, and they are not. A model that captures 90% of the SDF variance may produce a Fama-MacBeth slope of zero, or of the wrong sign, because the remaining 10% distributes across the cross-section in a pattern correlated with betas. Section 5 proposes three diagnostic practices, reporting $\rho_{\text{cs}}(\beta, \alpha)$, testing sensitivity to test assets, and separating model comparison from pricing inference, that would make the problem visible before researchers mistake it for a result.

A Including Characteristics as Controls

Suppose the cross-sectional regression includes a vector of characteristics c_i alongside betas:

$$r_{it} = \gamma_{0t} + \gamma_{1t}\,\hat{\beta}_i + \delta_t'\,c_i + e_{it}. \tag{51}$$

The same FWL logic from Section 2.5 applies: γ_{1t} is the return on a portfolio whose weights are determined by $\tilde{\beta}_i^{\perp c}$, the component of beta orthogonal to the characteristics. The expected slope is

$$E[\gamma_{1t}] = E[r_p] + \frac{\text{Cov}_{\text{cs}}(\beta^{\perp c},\, \alpha)}{\text{Var}_{\text{cs}}(\beta^{\perp c})}. \tag{52}$$

Including characteristics absorbs the OFB if and only if the residual betas, after removing the variation explained by characteristics, are uncorrelated with pricing errors. A sufficient condition is that the characteristics perfectly explain the cross-section of α_i, so $\mathrm{Cov}_{\mathrm{cs}}(\beta^{\perp c}, \alpha) = 0$ regardless of the distribution of $\beta^{\perp c}$. In this case, characteristics act as controls for the omitted SDF component by absorbing the pricing errors that would otherwise contaminate the slope.

This result connects to the debate between Daniel and Titman (1997) and Fama and French (1996) over whether characteristics or betas better explain expected returns. In the present framework, finding that characteristics "beat" betas in a horse-race regression is consistent with two interpretations. Under correct specification ($\alpha_i = 0$ for all i), characteristics have no explanatory power beyond betas; the finding implies misspecification. Under misspecification, characteristics may absorb the pricing errors, improving the slope estimate for the included factor. But characteristics can also absorb legitimate beta variation, reducing $\mathrm{Var}_{\mathrm{cs}}(\beta^{\perp c})$ and amplifying whatever residual OFB remains, exactly as adding correlated factors does in the multi-factor case. The net effect on the accuracy of γ_1 as an estimate of $E[r_p]$ is ambiguous without further structure.

The augmented regression (51) raises a structural problem. The pure beta regression $E[r_i] = \gamma_0 + \gamma_1 \beta_i$ is the pricing equation: it derives from the SDF, and γ_1 identifies the price of risk $E[r_p]$ under correct specification. The moment characteristics enter as additional regressors, this structural correspondence breaks. No standard SDF linear in traded factors generates $E[r_i] = \beta_i \lambda + \delta' c_i$ as a pricing equation, because characteristics are predetermined firm attributes, not covariances with marginal utility. The coefficient δ has a structural interpretation only under the auxiliary assumption that c_i proxies for loadings on omitted factors, in which case δ estimates the omitted prices of risk. But the cross-sectional regression alone cannot test this assumption: for any characteristic that predicts returns, one can always construct a factor (a long-short portfolio sorted on that characteristic) and rationalize the finding as a missing beta. The regression (51) is therefore a diagnostic for misspecification, not a pricing equation. It can reveal that the factor model is incomplete (if $\delta \neq 0$), but it cannot distinguish whether the incompleteness reflects missing risk factors, behavioral mispricing, or measurement error in betas, the last of which mechanically favors characteristics over estimated betas in any horse race.

References

Barillas, Francisco, Raymond Kan, Cesare Robotti, and Jay Shanken. 2020. "Model Comparison with Sharpe Ratios." *Journal of Financial and Quantitative Analysis* 55, no. 6 (September): 1840–1874. ISSN: 0022-1090, 1756-6916, accessed February 20, 2026. https://doi.org/10.1017/S0022109019000589.

Barillas, Francisco, and Jay Shanken. 2017. "Which Alpha?" *The Review of Financial Studies* 30, no. 4 (April 1, 2017): 1316–1338. ISSN: 0893-9454, accessed February 20, 2026. https://doi.org/10.1093/rfs/hhw101.

Black, Fischer. 1972. "Capital Market Equilibrium with Restricted Borrowing." *The Journal of Business* 45:444–455.

Bryzgalova, Svetlana, Alberto Quaini, Fabio Trojani, and Ming Yuan. 2023. "Tradable Factor Risk Premia and Oracle Tests of Asset Pricing Models." Pre-published, September 16, 2023. SSRN Scholarly Paper. Accessed March 4, 2026. https://doi.org/10.2139/ssrn.4574683. Social Science Research Network: 4574683.

Chen, Andrew Y., and Tom Zimmermann. 2022. "Open Source Cross-Sectional Asset Pricing." *Critical Finance Review* 11, no. 2 (May 3, 2022): 207–264. ISSN: 2164-5744. https://doi.org/10.1561/104.00000112.

Chen, Luyang, Markus Pelger, and Jason Zhu. 2024. "Deep Learning in Asset Pricing." *Management Science* 70, no. 2 (February): 714–750. ISSN: 0025-1909, accessed February 20, 2026. https://doi.org/10.1287/mnsc.2023.4695.

Chernov, Mikhail, Magnus Dahlquist, and Lars A. Lochstoer. 2025. "Unpriced Risks: Rethinking Cross-Sectional Asset Pricing." Pre-published, June 28, 2025. SSRN Scholarly Paper. Accessed February 20, 2026. https://doi.org/10.2139/ssrn.5341962. Social Science Research Network: 5341962.

Cochrane, John H. 2005. *Asset Pricing: Revised Edition.* Princeton University Press. ISBN: 978-1-4008-2913-2. https://books.google.co.uk/books?id=20pmeMaKNwsC.

Daniel, Kent, and Sheridan Titman. 1997. "Evidence on the Characteristics of Cross Sectional Variation in Stock Returns." *The Journal of Finance* 52 (1): 1–33. https://doi.org/10.1111/j.1540-6261.1997.tb03806.x.

Di Tella, Sebastian, Benjamin Hébert, Pablo Kurlat, and Qitong Wang. 2026. "The Zero-Beta Interest Rate." *Journal of Political Economy* (January 7, 2026). ISSN: 0022-3808, accessed January 23, 2026. https://doi.org/10.1086/740220.

Easterwood, Sara, and Bradley S. Paye. 2025. "High (on) Sharpe Ratios? Assessing the Performance of Data-Instigated Factor Models." Pre-published, December 15, 2025. SSRN Scholarly Paper. Accessed February 24, 2026. https://doi.org/10.2139/ssrn.4360788. Social Science Research Network: 4360788.

Fama, Eugene F., and Kenneth R. French. 1992. "The Cross-Section of Expected Stock Returns." *The Journal of Finance* 47, no. 2 (June 1, 1992): 427–465. ISSN: 15406261, accessed November 20, 2021. https://doi.org/10.1111/j.1540-6261.1992.tb04398.x.

———. 1993. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics* 33, no. 1 (February 1, 1993): 3–56. ISSN: 0304-405X, accessed November 20, 2021. https://doi.org/10.1016/0304-405X(93)90023-5.

———. 1996. "Multifactor Explanations of Asset Pricing Anomalies." *The journal of finance* 51 (1): 55–84.

Fama, Eugene F., and Kenneth R. French. 2015. "A Five-Factor Asset Pricing Model." *Journal of Financial Economics* 116, no. 1 (April 1, 2015): 1–22. ISSN: 0304-405X, accessed November 19, 2021. https://doi.org/http://dx.doi.org/10.1016/j.jfineco.2014.10.010.

———. 2018. "Choosing Factors." *Journal of Financial Economics* 128, no. 2 (May 1, 2018): 234–252. ISSN: 0304-405X, accessed February 20, 2026. https://doi.org/10.1016/j.jfineco.2018.02.012.

Fama, Eugene F., and James D. MacBeth. 1973. "Risk, Return, and Equilibrium: Empirical Tests." *Journal of Political Economy* 81 (3): 607–636. ISSN: 00223808, 1537534X. JSTOR: 1831028. http://www.jstor.org/stable/1831028.

Frazzini, Andrea, and Lasse Heje Pedersen. 2014. "Betting against Beta." *Journal of Financial Economics* 111:1–25.

Giannetti, Antoine. 2024. "A Simple Test of Misspecification for Linear Asset Pricing Models." *Financial Markets and Portfolio Management* 38, no. 3 (September 1, 2024): 305–330. ISSN: 2373-8529, accessed February 24, 2026. https://doi.org/10.1007/s11408-024-00445-6.

Gibbons, Michael R, Stephen A Ross, and Jay Shanken. 1989. "A Test of the Efficiency of a Given Portfolio." *Econometrica* 57 (5): 1121–1152. ISSN: 00129682, 14680262. JSTOR: 1913625. http://www.jstor.org/stable/1913625.

Giglio, Stefano, and Dacheng Xiu. 2021. "Asset Pricing with Omitted Factors." *Journal of Political Economy* 129 (7): 1947–1990. ISSN: 1537534X. https://doi.org/10.1086/714090.

Giglio, Stefano, Dacheng Xiu, and Dake Zhang. 2021. "Test Assets and Weak Factors." *SSRN Electronic Journal,* https://doi.org/10.2139/ssrn.3768081.

Gospodinov, Nikolay, Raymond Kan, and Cesare Robotti. 2014. “Misspecification-Robust Inference in Linear Asset-Pricing Models with Irrelevant Risk Factors.” *The Review of Financial Studies* 27, no. 7 (July 1, 2014): 2139–2170. ISSN: 0893-9454, accessed February 20, 2026. https://doi.org/10.1093/rfs/hht135.

Hansen, Lars Peter, and Ravi Jagannathan. 1991. “Implications of Security Market Data for Models of Dynamic Economies.” *Journal of Political Economy* 99 (2): 225–262. ISSN: 0022-3808. https://doi.org/10.1086/261749.

Hansen, Lars Peter, and Scott F. Richard. 1987. “The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models.” *Econometrica* 55 (3): 587–613. ISSN: 0012-9682, accessed February 20, 2026. https://doi.org/10.2307/1913601. JSTOR: 1913601.

Harvey, Campbell R, Yan Liu, and Heqing Zhu. 2016. “. . . and the Cross-Section of Expected Returns.” *The Review of Financial Studies* 29 (1): 5–68. https://doi.org/10.1093/rfs/hhv059.

Hollstein, Fabian, and Marcel Prokopczuk. 2022. “Testing Factor Models in the Cross-Section.” *Journal of Banking & Finance* 145 (December 1, 2022): 106626. ISSN: 0378-4266, accessed February 24, 2026. https://doi.org/10.1016/j.jbankfin.2022.106626.

Kan, Raymond, Cesare Robotti, and Jay Shanken. 2013. “Pricing Model Performance and the Two-Pass Cross-Sectional Regression Methodology.” *The Journal of Finance* 68 (6): 2617–2649. ISSN: 1540-6261, accessed February 20, 2026. https://doi.org/10.1111/jofi.12035.

Kan, Raymond, and Chu Zhang. 1999. “Two-Pass Tests of Asset Pricing Models with Useless Factors.” *The Journal of Finance* 54 (1): 203–235. ISSN: 1540-6261, accessed February 27, 2026. https://doi.org/10.1111/0022-1082.00102.

Kandel, Shmuel, and Robert F. Stambaugh. 1995. “Portfolio Inefficiency and the Cross-section of Expected Returns.” *The Journal of Finance* 50 (1): 157–184. ISSN: 1540-6261, accessed February 20, 2026. https://doi.org/10.1111/j.1540-6261.1995.tb05170.x.

Kelly, Bryan T, Seth Pruitt, and Yinan Su. 2019. “Characteristics Are Covariances: A Unified Model of Risk and Return.” *Journal of Financial Economics* 134 (3): 501–524. ISSN: 0304-405X. https://doi.org/10.1016/j.jfineco.2019.05.001.

Kozak, Serhiy, and Stefan Nagel. 2024. “When Do Cross-Sectional Asset Pricing Factors Span the Stochastic Discount Factor?” *SSRN Electronic Journal* (May 29, 2024). Accessed November 18, 2024. https://doi.org/10.2139/SSRN.4443643.

Kozak, Serhiy, Stefan Nagel, and Shrihari Santosh. 2018. “Interpreting Factor Models.” *The Journal of Finance* 73 (3): 1183–1223. ISSN: 15406261. https://doi.org/10.1111/jofi.12612.

———. 2020. “Shrinking the Cross-Section.” *Journal of Financial Economics* 135, no. 2 (February 1, 2020): 271–292. ISSN: 0304-405X, accessed February 20, 2026. https://doi.org/10.1016/j.jfineco.2019.06.008.

Lettau, Martin, and Markus Pelger. 2020. “Factors That Fit the Time Series and Cross-Section of Stock Returns.” Edited by Stijn Van Nieuwerburgh. *The Review of Financial Studies* 33 (5): 2274–2325. ISSN: 0893-9454. https://doi.org/10.1093/rfs/hhaa020.

Lewellen, Jonathan, Stefan Nagel, and Jay Shanken. 2010. “A Skeptical Appraisal of Asset Pricing Tests.” *Journal of Financial Economics* 96 (2): 175–194. ISSN: 0304-405X. https://doi.org/10.1016/j.jfineco.2009.09.001.

McLean, R David, and Jeffrey Pontiff. 2016. “Does Academic Research Destroy Stock Return Predictability?” *The Journal of Finance* 71 (1): 5–32. https://doi.org/10.1111/jofi.12365.

Roll, Richard. 1977. "A Critique of the Asset Pricing Theory's Tests Part I: On Past and Potential Testability of the Theory." *Journal of Financial Economics* 4 (2): 129–176. ISSN: 0304-405X. https://doi.org/10.1016/0304-405X(77)90009-5.

Shanken, Jay. 1987. "Multivariate Proxies and Asset Pricing Relations: Living with the Roll Critique." *Journal of Financial Economics* 18, no. 1 (March 1, 1987): 91–110. ISSN: 0304-405X, accessed February 24, 2026. https://doi.org/10.1016/0304-405X(87)90062-6.

———. 1992. "On the Estimation of Beta-Pricing Models." *The review of financial studies* 5 (1): 1–33.

Wang, Rong. 2024. "Resolving the Zero-Beta Rate Puzzle." Pre-published, July 27, 2024. SSRN Scholarly Paper. Accessed March 4, 2026. https://doi.org/10.2139/ssrn.4942202. Social Science Research Network: 4942202.