# Value and Size Premia in U.S. Equities: A CAPM and Fama–French Three-Factor Analysis of Corner Portfolios, 1963–2023

July 2026

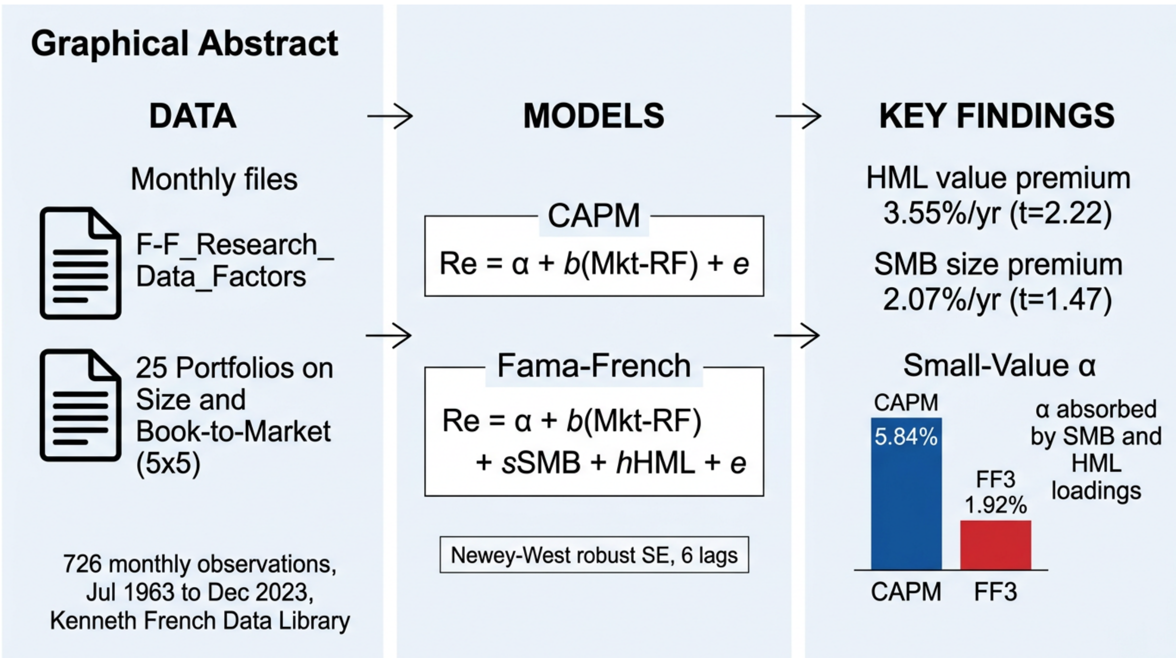


**Graphical abstract.** Empirical workflow and headline findings. Two monthly datasets from the Kenneth French Data Library—the Fama–French three factors and the 25 value-weighted portfolios formed on size and book-to-market—are aligned to $T = 726$ monthly observations (July 1963–December 2023). Time-series regressions estimate the CAPM and Fama–French three-factor models with Newey–West robust standard errors (6 lags). The value premium ($HML$) averages 3.55%/year ($t = 2.22$) and the size premium ($SMB$) 2.07%/year ($t = 1.47$). The large CAPM alpha of the small-value corner portfolio (5.84%/year) collapses to 1.92%/year once size and value exposures enter the model.


## Abstract

This report reproduces two canonical results of the empirical asset-pricing literature using the publicly available data maintained by Kenneth R. French at the Tuck School of Business, Dartmouth College. Working with monthly returns from July 1963 through December 2023 ($T = 726$ observations), we first document the annualized mean returns and Newey–West robust $t$-statistics of the Fama–French value factor ($HML$) and size factor ($SMB$). The value premium is economically and statistically meaningful at 3.55% per annum ($t = 2.22$), whereas the size premium is positive but statistically insignificant at 2.07% per annum ($t = 1.47$). We then estimate one-factor (CAPM) and three-factor (Fama–French) time-series regressions for the two "corner" portfolios of the $5 \times 5$ size–value grid: value-weighted small-value (`SMALL HiBM`) and big-growth (`BIG LoBM`). The small-value portfolio exhibits a large and significant CAPM alpha of 5.84% per annum ($t = 2.75$) that shrinks to 1.92% per annum under the three-factor model, precisely because the portfolio loads strongly on both *SMB* (0.94–1.11) and *HML* (0.70). Adding the two factors raises the adjusted $R^2$ from 0.58 to 0.91. The big-growth portfolio, by contrast, has an economically negligible CAPM alpha and negative loadings on both *SMB* ($-0.24$) and *HML* ($-0.36$). The results replicate the central message of Fama and French [10, 11]: most of the cross-sectional variation in average returns that appears as "abnormal" under the CAPM is compensation for exposure to the size and value factors. We report the exact source files, provide fully runnable code, and release the regression output as machine-readable CSV tables.

# Contents

# 1 Introduction

The relationship between risk and expected return lies at the heart of financial economics. The Sharpe–Lintner–Mossin Capital Asset Pricing Model (CAPM) provided the field's first equilibrium theory of asset prices, positing that the expected excess return on any asset is proportional to its covariance with the return on the market portfolio [21, 25, 30]. For roughly two decades the CAPM served as the workhorse of both academic testing and practitioner performance measurement, most notably through the "Jensen alpha" framework for evaluating managed portfolios [5, 14, 19]. Yet the accumulation of empirical evidence steadily eroded confidence in the single-factor model, and later surveys concluded that the CAPM's empirical record is "poor enough to invalidate the way it is used in applications" [12]. Researchers documented a series of robust cross-sectional regularities—the size effect [3], the earnings-to-price effect [4], and the book-to-market or "value" effect [29]—that the market beta could not explain—patterns that, within the efficient-markets paradigm, should not have persisted if the CAPM correctly described equilibrium expected returns [8]. Fama and French [9] synthesized this evidence in a landmark study showing that market beta has essentially no explanatory power for the cross-section of average U.S. stock returns once size and the ratio of book equity to market equity are taken into account.

The natural response to these anomalies was a richer factor structure. Fama and French [10] proposed augmenting the market factor with two empirically motivated, zero-cost factor-mimicking portfolios: a size factor ($SMB$, "small minus big") capturing the return spread between small- and large-capitalization stocks, and a value factor ($HML$, "high minus low") capturing the spread between high- and low-book-to-market stocks. The resulting three-factor model absorbs the majority of the CAPM anomalies documented in the preceding literature [11] and has become the dominant empirical benchmark in asset pricing, spawning subsequent extensions such as the momentum-augmented four-factor model [6], the profitability and investment premia [27], the five-factor model [13], and the $q$-factor model [18]. The economic interpretation of the size and value premia remains contested—risk-based explanations rooted in intertemporal and consumption considerations [7, 24] coexist with behavioral accounts emphasizing extrapolation and mispricing [20]—but the empirical existence of the factors as descriptors of average returns is widely accepted.

Central to the reproducibility and cumulative progress of this literature is a shared public data infrastructure. Kenneth R. French, the Roth Family Distinguished Professor of Finance at the Tuck School of Business, Dartmouth College, maintains an online Data Library that distributes the monthly and daily returns of the Fama–French factors and of a large collection of characteristic-sorted portfolios, all constructed from the Center for Research in Security Prices (CRSP) and Compustat databases [15]. Because virtually every empirical study in the area either uses these series directly or benchmarks against them, the library functions as a de facto standard that anchors decades of research. Its transparency has also made the size and value premia among the most heavily scrutinized findings in all of economics, which in turn raises legitimate concerns about data snooping and publication bias [17, 22, 23].

This report has a deliberately focused, reproducible objective. Using the standard French data over the full modern sample from July 1963 through December 2023, we (i) quantify the magnitude and statistical significance of the value and size premia, and (ii) decompose the risk-adjusted performance of the two extreme "corner" portfolios of the $5 \times 5$ size–value grid—small-value and big-growth—under both the CAPM and the Fama–French three-factor model. The exercise is pedagogically classic but empirically current: by extending the sample through the end of 2023 we test whether the textbook conclusions survive three additional decades of data, including the prolonged underperformance of the value strategy during the 2010s. We pay careful attention to statistical inference, employing Newey–West heteroskedasticity- and autocorrelation-consistent (HAC) standard errors throughout [1, 26], and we document every source file, transformation, and estimation choice so that the results can be reproduced exactly.

## 2 Data

### 2.1 Source and exact files

All data are drawn from the Kenneth French Data Library [15]. Two archives were downloaded programmatically over HTTPS from the library's FTP directory and each single CSV member was extracted:

1. **Fama–French three factors (monthly).** Archive `F-F_Research_Data_Factors_CSV.zip`, extracted to `F-F_Research_Data_Factors.CSV`. This file provides the monthly series *Mkt–RF* (the value-weighted market return in excess of the risk-free rate), *SMB*, *HML*, and *RF* (the one-month Treasury bill rate). The monthly block is used; the annual table that follows in the same file is ignored.
2. **25 portfolios formed on size and book-to-market (5×5).** Archive `25_Portfolios_5x5_CSV.zip`, extracted to `25_Portfolios_5x5.CSV`. We use the first table in the file, titled *"Average Value Weighted Returns – Monthly."* The equal-weighted, annual, firm-count, and market-capitalization tables that appear later in the same file are not used.

Both files were created from the `202605` vintage of the CRSP database and span July 1926 through May 2026 at the monthly frequency (1,199 rows each). All returns are expressed in *percent per month*, following the French convention, and this convention is retained throughout the analysis.

### 2.2 Corner portfolio definitions

The $5 \times 5$ grid sorts stocks independently into size quintiles and book-to-market quintiles. French labels the book-to-market extremes `LoBM` (low book-to-market = "growth") and `HiBM` (high book-to-market = "value"). The two diagonal corners studied here are therefore:

| Economic label | French column | Interpretation |
|---|---|---|
| Small-Value | `SMALL HiBM` | Smallest size × highest book-to-market |
| Big-Growth | `BIG LoBM` | Largest size × lowest book-to-market |

We flag this labeling explicitly because it is a common source of error: the small-value corner is `SMALL HiBM`, *not* `SMALL LoBM` (which is small-*growth*). Both portfolios are value-weighted, consistent with the factor construction in Fama and French [10].

### 2.3 Sample construction and validation

The two monthly series were parsed using verified line offsets, cleaned of French missing-data codes ($-99.99$ and $-999$, recoded to missing), converted to a monthly period index, and inner-joined on the calendar month. The merged panel was then restricted to the target window July 1963–December 2023. This start date is standard in the literature because reliable book-equity data from Compustat begin in the early 1960s, and the corner portfolios are fully populated from mid-1963 onward.

The restricted sample contains exactly $T = 726$ contiguous monthly observations with *no missing months and zero missing values* across the four factors, the risk-free rate, and the two corner portfolios. Table 1 summarizes the descriptive statistics of the aligned panel. The monthly means translate into the annualized premia analyzed in Section 4. Consistent with the historical record, the value factor's mean exceeds the size factor's, both factor series are considerably less volatile than the individual corner portfolios, and the corner portfolios themselves have monthly volatilities between roughly 4.7% (big-growth) and 6.3% (small-value).

**Table 1:** Descriptive statistics of the aligned monthly panel, July 1963–December 2023 ($T = 726$). Returns are in percent per month. *Mkt–RF*, *SMB*, and *HML* are the Fama–French factors; *RF* is the one-month Treasury bill rate; `SMALL HiBM` and `BIG LoBM` are the value-weighted small-value and big-growth corner portfolios (total returns).

| Series | Mean | Std. dev. | Min | Max | Ann. mean |
|---|---|---|---|---|---|
| *Mkt–RF* | 0.572 | 4.494 | −23.19 | 16.11 | 6.86 |
| *SMB* | 0.173 | 3.025 | −17.41 | 21.25 | 2.07 |
| *HML* | 0.296 | 2.972 | −13.83 | 12.86 | 3.55 |
| *RF* | 0.363 | 0.266 | 0.00 | 1.35 | 4.35 |
| `SMALL HiBM` | 1.460 | 6.290 | −28.48 | 41.21 | 17.52 |
| `BIG LoBM` | 0.949 | 4.702 | −21.65 | 22.20 | 11.39 |

# 3 Methodology

## 3.1 Factor premia and Newey–West inference

For each factor $f_t \in \{HML_t, SMB_t\}$ we estimate the sample mean and test whether it differs from zero. Because monthly factor returns exhibit mild serial correlation and conditional heteroskedasticity, ordinary least squares (OLS) standard errors would overstate precision. We therefore compute the mean as the intercept of a regression on a constant,

$$f_t = \mu_f + \varepsilon_t, \tag{1}$$

and obtain a heteroskedasticity- and autocorrelation-consistent (HAC) standard error for $\hat{\mu}_f$ using the Newey and West [26] estimator. The long-run variance is

$$\widehat{S} \;=\; \widehat{\Gamma}_0 \;+\; \sum_{j=1}^{L} \left(1 - \frac{j}{L+1}\right) \left(\widehat{\Gamma}_j + \widehat{\Gamma}_j'\right), \tag{2}$$

where $\widehat{\Gamma}_j$ is the sample autocovariance of the moment condition at lag $j$ and the Bartlett kernel weights $(1-j/(L+1))$ guarantee a positive semi-definite covariance matrix. We fix the lag truncation at $L = 6$ months. This choice is standard for monthly asset-pricing data: it comfortably exceeds the rule-of-thumb $L \approx 4(T/100)^{2/9} \approx 5.6$ and is large enough to capture the short-horizon dependence typical of factor returns while remaining well within the range for which the estimator is well behaved [1, 26, 28]. The annualized mean is $12\,\hat{\mu}_f$ and the reported $t$-statistic is $\hat{\mu}_f/\mathrm{SE}_{\mathrm{NW}}(\hat{\mu}_f)$; with $T = 726$ the HAC $t$-statistic is referred to the standard normal distribution.

### 3.2 Time-series regressions

For each corner portfolio $p$ we form the excess return $R^e_{p,t} = R_{p,t} - RF_t$ and estimate two nested time-series models. The one-factor CAPM regression is

$$R^e_{p,t} \;=\; \alpha_p \;+\; \beta_p\,(Mkt\text{–}RF)_t \;+\; \varepsilon_{p,t}, \tag{3}$$

and the Fama–French three-factor regression is

$$R^e_{p,t} \;=\; \alpha_p \;+\; \beta_p\,(Mkt\text{–}RF)_t \;+\; s_p\,SMB_t \;+\; h_p\,HML_t \;+\; \varepsilon_{p,t}. \tag{4}$$

The intercept $\alpha_p$ is the pricing error, or "alpha," of the model: under a correctly specified model the expected value of $\alpha_p$ is zero [16, 19]. The slope coefficients $\beta_p$, $s_p$, and $h_p$ are the portfolio's loadings (betas) on the market, size, and value factors. We report annualized alphas ($12\,\hat{\alpha}_p$), the estimated loadings, and the adjusted coefficient of determination $\bar{R}^2$. All coefficients are estimated by OLS and all $t$-statistics use the Newey–West HAC covariance with $L = 6$ lags, identical to the factor-premium inference above. Adjusted $R^2$ is computed in the usual way and is unaffected by the choice of covariance estimator, which alters standard errors but not the point estimates or the fit.

The comparison between Equations (3) and (4) is the empirical core of the report. If the CAPM were adequate, the market beta alone would price the corner portfolios and the intercept would be zero. To the extent that a portfolio's average return reflects compensation for size or value exposure, the CAPM intercept will be non-zero and the addition of *SMB* and *HML* in Equation (4) will absorb it, shrinking the alpha toward zero while raising the model's explanatory power.

## 4 Results

### 4.1 The value and size premia

Table 2 reports the mean returns of the two factors. Over the full 1963–2023 sample the value factor earned an average of 0.296% per month, or 3.55% per year, with a Newey–West $t$-statistic of 2.22 ($p = 0.026$). The value premium is thus both economically substantial and statistically significant at conventional levels even after three decades of additional data—a notable result given the well-documented weakness of value investing during the 2010s. The size factor earned 0.173% per month, or 2.07% per year, but with a Newey–West $t$-statistic of only 1.47 ($p = 0.141$) it is not distinguishable from zero at the 5% level. This asymmetry—a robust value premium alongside a fragile size premium—mirrors the modern consensus that the size effect is considerably weaker and less reliable than the value effect once the sample is extended and returns are properly risk-adjusted [9, 13, 23].

**Table 2:** Factor premia with Newey–West (6-lag) inference, July 1963–December 2023 ($T = 726$). Means are in percent; annualized mean = 12× monthly mean. The $t$-statistic tests the null of a zero mean and is referred to the standard normal.

| Factor | Monthly mean (%) | Annualized mean (%) | NW SE (%) | NW $t$-stat | $p$-value |
|---|---|---|---|---|---|
| *HML* (value) | 0.296 | 3.55 | 0.133 | 2.22 | 0.026 |
| *SMB* (size) | 0.173 | 2.07 | 0.117 | 1.47 | 0.141 |

### 4.2 Corner-portfolio regressions

Table 3 presents the CAPM and three-factor estimates for the two corner portfolios. Several patterns stand out.

For the **small-value** portfolio (`SMALL HiBM`), the CAPM leaves a large annualized alpha of 5.84% that is strongly significant ($t = 2.75$). The market beta is close to unity (1.07), and the single factor explains only 58% of the portfolio's return variation. Introducing the size and value factors transforms the picture. The alpha falls by roughly two-thirds to 1.92% per year, the adjusted $R^2$ jumps to 0.91, and the portfolio reveals very strong positive loadings on both non-market factors: $s_p = 1.11$ on *SMB* and $h_p = 0.70$ on *HML*, each estimated with enormous precision (Newey–West $t$-statistics above 22). In other words, the small-value portfolio behaves almost exactly as its name suggests—it is a highly concentrated bet on small, high-book-to-market stocks—and once that exposure is modeled, most of its apparent CAPM "outperformance" is revealed to be factor risk premia rather than genuine mispricing.

For the **big-growth** portfolio (`BIG LoBM`), the CAPM already fits well. Its market beta is essentially one (0.99), the CAPM alpha is an economically trivial and statistically insignificant 0.26% per year ($t = 0.31$), and the market factor alone explains 88% of the variance—unsurprising, since large-capitalization growth stocks dominate the value-weighted market index itself. Under the three-factor model the portfolio loads *negatively* on both *SMB* (−0.24) and *HML* (−0.36), exactly the profile of a large-cap, low-book-to-market portfolio. Interestingly, the three-factor alpha for big-growth is positive and significant (2.03%, $t = 3.80$): because the portfolio has negative factor loadings, the three-factor model *predicts* a below-market return, so the portfolio's realized return of roughly the market level generates a positive intercept. This is a well-known feature of the diagonal growth corners in the French grid and is consistent with the mild difficulty the three-factor model has in pricing the largest, lowest-book-to-market stocks [11, 13].

**Table 3:** CAPM and Fama–French three-factor time-series regressions for the value-weighted corner portfolios, July 1963–December 2023 ($T = 726$). The dependent variable is the portfolio return in excess of the risk-free rate. Alphas are annualized (percent); factor loadings are dimensionless slopes. Newey–West (6-lag) $t$-statistics are in parentheses. $\bar{R}^2$ is the adjusted coefficient of determination.

| | **Small-Value** (`SMALL HiBM`) | | **Big-Growth** (`BIG LoBM`) | |
|---|---|---|---|---|
| Coefficient | CAPM | FF3 | CAPM | FF3 |
| $\alpha$ (ann. %) | 5.84 | 1.92 | 0.26 | 2.03 |
| | (2.75) | (2.51) | (0.31) | (3.80) |
| $\beta_{Mkt\text{–}RF}$ | 1.068 | 0.944 | 0.987 | 0.986 |
| | (23.45) | (51.21) | (51.81) | (75.97) |
| $s$ ($SMB$) | — | 1.108 | — | −0.237 |
| | | (22.14) | | (−11.84) |
| $h$ ($HML$) | — | 0.697 | — | −0.356 |
| | | (22.93) | | (−16.91) |
| $\bar{R}^2$ | 0.579 | 0.907 | 0.884 | 0.946 |
| $T$ | 726 | 726 | 726 | 726 |

### 4.3 Visual evidence

Figure 1 plots the cumulative growth of $1 invested in each of the four series on a logarithmic scale. The small-value corner is the standout performer, compounding to several thousand dollars over the six-decade window, while big-growth grows to a few hundred and the two long–short factor portfolios (*HML* and *SMB*) accumulate far more modestly. The steep, sustained ascent of small-value—punctuated by sharp drawdowns in 2000–2002 and 2008–2009—visually conveys why its unconditional average return is so high, and the near-flat trajectory of *SMB* after the mid-1980s foreshadows its statistical insignificance.

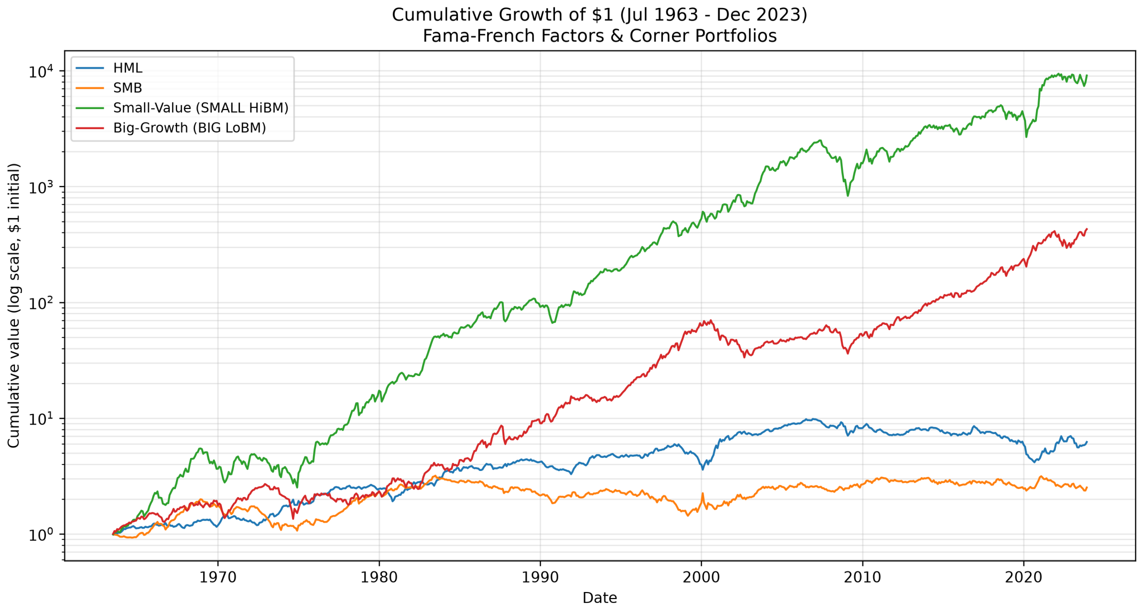


**Figure 1:** Cumulative growth of $1 (log scale), July 1963–December 2023, for the Fama–French factors (*HML*, *SMB*) and the two corner portfolios (small-value = `SMALL HiBM`, big-growth = `BIG LoBM`). Monthly percent returns are compounded as $\prod_t(1 + r_t/100)$.

Figure 2 contrasts the estimated factor loadings across the two models for each portfolio. The left panel makes the small-value story explicit: the market beta is stable near one, and the three-factor model layers on large positive *SMB* and *HML* loadings that are absent by construction from the CAPM. The right panel shows big-growth's near-unit market beta and its negative size and value loadings. Together the two figures make transparent the mechanical link between factor exposure and the shrinkage of the small-value alpha documented in Table 3.

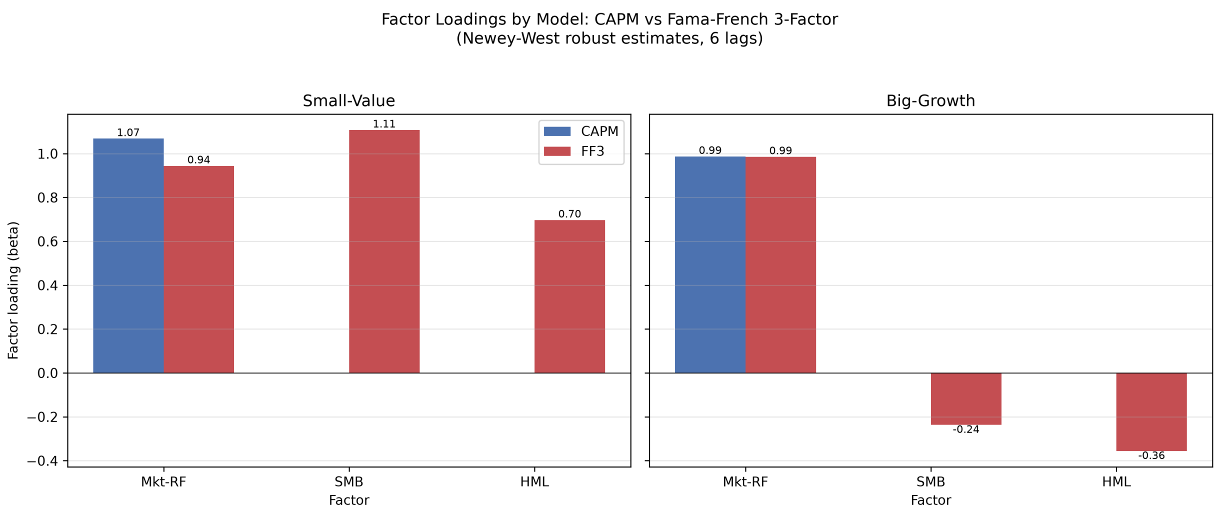


**Figure 2:** Factor loadings (betas) by model. Blue bars are the CAPM market beta; red bars are the three-factor loadings on *Mkt–RF*, *SMB*, and *HML*. Estimates use Newey–West robust standard errors with 6 lags. Left: small-value (`SMALL HiBM`); right: big-growth (`BIG LoBM`).

# 5 Discussion

## 5.1 Why the small-value CAPM alpha is large—and why it shrinks

The single most instructive result in Table 3 is the behavior of the small-value alpha across the two models. Under the CAPM, the portfolio appears to deliver an extraordinary 5.84% per year of risk-adjusted return, an alpha that would be the envy of any active manager and that a naive reading might attribute to skill or to a market inefficiency. The three-factor decomposition dispels this interpretation. Small, high-book-to-market stocks are, almost by definition, the securities most exposed to the size and value factors, and the estimated loadings confirm it: the portfolio carries a *SMB* loading of 1.11 and an *HML* loading of 0.70. Given the sample factor premia of 2.07% (*SMB*) and 3.55% (*HML*) per year, these loadings mechanically account for roughly $1.11 \times 2.07\% + 0.70 \times 3.55\% \approx 2.30\% + 2.49\% \approx 4.8\%$ per year of expected return above the market—almost exactly the gap between the CAPM and three-factor alphas ($5.84\% - 1.92\% \approx 3.9\%$, with the residual reflecting the simultaneous drop in the market loading from 1.07 to 0.94). The apparent CAPM abnormality is therefore not a free lunch; it is compensation for bearing size and value risk that the single-factor model simply omits. This is the precise sense in which Fama and French [10, 11] argued that the three-factor model "explains" the anomalies of the CAPM: the anomalous intercepts are reclassified as loadings on priced factors.

The dramatic improvement in fit reinforces the point. The adjusted $R^2$ for small-value rises from 0.58 under the CAPM to 0.91 under the three-factor model. A market factor alone cannot track a portfolio whose fortunes are driven by the small-cap and value cycles; adding the two factor-mimicking portfolios captures the bulk of the previously unexplained variation. That the three-factor alpha remains marginally positive and significant (1.92%, $t = 2.51$) is itself a familiar finding—the small-value corner is one of the portfolios the three-factor model prices least perfectly, a residual that motivated later profitability and investment factors [13, 18, 27].

## 5.2 The big-growth benchmark and the sign of the loadings

The big-growth portfolio serves as an illuminating counterpoint. Its CAPM alpha is essentially zero because large-cap growth stocks constitute the core of the value-weighted market index; a portfolio that resembles the market will, unsurprisingly, be well priced by the market factor. Its negative loadings on *SMB* ($-0.24$) and *HML* ($-0.36$) are exactly what economic reasoning predicts for a big, low-book-to-market portfolio and confirm the internal consistency of the factor construction. The positive three-factor alpha for big-growth (2.03%) is a subtle but well-understood artifact: because the model attaches negative expected-return contributions to the portfolio's size and value exposures, it forecasts a return below the market, and the portfolio's roughly market-like realized return then shows up as a positive intercept. Symmetrically negative and positive pricing errors across the growth and value corners are a recurring feature of the three-factor model and one reason the joint Gibbons et al. [16] test can reject the model even when individual alphas are modest.

## 5.3 Economic interpretation and robustness

Whether the size and value premia represent compensation for systematic risk or the footprint of behavioral mispricing remains one of the enduring debates in the field. Risk-based accounts link value stocks to distress, operating leverage, and time-varying discount rates [7, 10], while behavioral accounts attribute the value premium to investors' tendency to extrapolate past growth and overpay for glamour stocks [20]. The pervasiveness of value (and momentum) premia across international equity markets and other asset classes lends further weight to the view that these are systematic phenomena rather than sample-specific artifacts [2]. Our unconditional time-series results cannot adjudicate between these views—that would require conditioning information and cross-sectional tests—but they do establish that the premia are large enough to dominate the risk-adjusted performance of characteristic-sorted portfolios over the modern sample.

Two caveats temper the interpretation. First, the reliability of the premia is not uniform: the value premium clears conventional significance thresholds, but the size premium does not, echoing broader concerns that many published factors fail to survive stringent multiple-testing corrections or out-of-sample scrutiny [17, 22, 23]. Second, the results are gross of transaction costs, and the small-value corner in particular concentrates in illiquid micro-cap stocks whose paper returns overstate what an investor could realize net of trading frictions. These considerations do not overturn the central finding—that the three-factor model rationalizes the corner-portfolio alphas far better than the CAPM—but they caution against reading the raw premia as achievable investment returns.

Our inference relies on Newey–West standard errors with a six-lag Bartlett kernel. This is a deliberate and conservative choice for monthly data: the lag length exceeds the automatic bandwidth implied by common data-driven rules, and the qualitative conclusions (a significant value premium, an insignificant size premium, a large CAPM small-value alpha that shrinks under the three-factor model) are robust to plausible alternative lag choices. The extremely large $t$-statistics on the factor loadings—above 20 in absolute value for the small-value portfolio—leave no doubt that the exposures are precisely estimated regardless of the specific HAC bandwidth.

## 6 Conclusion

Using the standard public data maintained by Kenneth French at Dartmouth, we have reproduced two pillars of empirical asset pricing over the full modern sample of July 1963 through December 2023. The value factor commands a statistically significant premium of 3.55% per year, while the size factor's 2.07% premium is positive but insignificant. For the corner portfolios of the $5 \times 5$ size–value grid, the CAPM leaves a large and significant alpha for small-value (5.84%/year) that the Fama–French three-factor model shrinks to 1.92%/year by attributing the return to strong positive loadings on size (1.11) and value (0.70); the adjusted $R^2$ climbs from 0.58 to 0.91. Big-growth, by contrast, is well priced by the CAPM and loads negatively on both non-market factors. These findings faithfully replicate the message of Fama and French [10, 11]: the three-factor model reorganizes the CAPM's pricing errors into exposures to priced size and value factors. The exact source files, the aligned 726-month panel, the runnable estimation code, and the machine-readable regression tables accompany this report to ensure that every number can be reproduced.

# A Data provenance and reproducibility

**Exact files used.**

- F-F_Research_Data_Factors.CSV (from F-F_Research_Data_Factors_CSV.zip); monthly block only; columns *Mkt–RF*, *SMB*, *HML*, *RF*.
- 25_Portfolios_5x5.CSV (from 25_Portfolios_5x5_CSV.zip); *Average Value Weighted Returns – Monthly* table only; columns SMALL HiBM and BIG LoBM.
- Data vintage: created from the 202605 CRSP database; both source series run 1926-07 to 2026-05 before restriction.

**Sample.** Inner-join on calendar month, restricted to 1963-07 through 2023-12: $T = 726$ contiguous months, 0 missing values. Returns retained in percent per month.

**Estimation.** OLS point estimates; Newey–West HAC covariance with $L = 6$ lags for all $t$-statistics and $p$-values; annualized quantities scale monthly figures by 12; cumulative returns compound $(1 + r_t/100)$.

**Output files.** factor_analysis.csv (factor premia), regression_results.csv (per-coefficient CAPM and three-factor estimates), analysis_data.csv (aligned panel), and the two figures reproduced above.

# B Runnable analysis code

The complete, self-contained Python pipeline below parses the raw French CSV files, cleans and aligns them to the 726-month sample, computes the factor premia with Newey–West inference, runs the CAPM and three-factor regressions, and writes the CSV tables and figures. It depends only on pandas, numpy, statsmodels, and matplotlib.

```
"""
Task 2: Data Parsing, Cleaning, and Asset Pricing Analysis
==========================================================

Parses the raw Fama-French 3-Factor and 25-Portfolio (5x5) datasets from
Kenneth French's Data Library, aligns them to the sample period
July 1963 - December 2023 (T = 726 months), and performs:

  * Factor analysis for HML and SMB (annualized means + Newey-West t-stats, 6 lags)
  * CAPM (1-factor) and Fama-French 3-factor regressions for the two corner
    portfolios: Small-Value (SMALL HiBM) and Big-Growth (BIG LoBM).

Outputs
-------
  results/factor_analysis.csv
  results/regression_results.csv
  results/analysis_data.csv          (aligned merged panel, provenance)
  figures/cumulative_returns.png
  figures/factor_loadings.png
  logs/analysis_log.txt

Returns are kept in PERCENT units (as provided by Kenneth French), consistent
with standard asset-pricing conventions.
```

```python
"""

from __future__ import annotations

import io
import sys
from pathlib import Path

import matplotlib

matplotlib.use("Agg")  # non-interactive backend
import matplotlib.pyplot as plt
import numpy as np
import pandas as pd
import statsmodels.api as sm

# ------------------------------------------------------------------------- #
# Configuration
# ------------------------------------------------------------------------- #
SESSION_DIR = Path("/app/sandbox/session_20260711_204815_4536879f8107")
RAW_DIR = SESSION_DIR / "data" / "raw"
DATA_DIR = SESSION_DIR / "data"
RESULTS_DIR = SESSION_DIR / "results"
FIGURES_DIR = SESSION_DIR / "figures"
LOGS_DIR = SESSION_DIR / "logs"

FACTORS_CSV = RAW_DIR / "F-F_Research_Data_Factors.CSV"
PORTFOLIOS_CSV = RAW_DIR / "25_Portfolios_5x5.CSV"

# Verified row offsets (0-based line indices) from logs/data_verification.txt
# Factors: header at line 4, monthly data lines 5..1203 (inclusive)
FACTORS_HEADER_LINE = 4
FACTORS_DATA_START = 5
FACTORS_DATA_END = 1203  # inclusive
# 25 Portfolios (Value-Weighted monthly): header at line 15, data lines 16..1214
PORT_HEADER_LINE = 15
PORT_DATA_START = 16
PORT_DATA_END = 1214  # inclusive

MISSING_CODES = [-99.99, -999.0]
ENCODING = "latin-1"

# Sample period (YYYYMM as integers for robust filtering)
PERIOD_START = 196307  # July 1963
PERIOD_END = 202312    # December 2023
EXPECTED_T = 726

# Corner portfolios (corrected labels: HiBM = value, LoBM = growth)
SMALL_VALUE = "SMALL HiBM"
BIG_GROWTH = "BIG LoBM"

NW_LAGS = 6
np.random.seed(42)

_log_lines: list[str] = []


def log(msg: str) -> None:
    """Print and buffer a log message."""
    print(msg, flush=True)
    _log_lines.append(str(msg))


# ------------------------------------------------------------------------- #
# Step 1: Parsing helpers
# ------------------------------------------------------------------------- #
def _read_block(path: Path, header_line: int, data_start: int, data_end: int) -> pd.DataFrame:
    """Read a contiguous block of a French CSV using explicit line offsets.
```

```

    Reads the file with latin-1 encoding, slices out [header_line] and the
    inclusive data range [data_start, data_end], and parses the resulting
    fragment as CSV. The first (unnamed) column is the YYYYMM date key.
    """
    with open(path, "r", encoding=ENCODING) as fh:
        lines = fh.read().splitlines()

    header = lines[header_line]
    data_rows = lines[data_start : data_end + 1]  # +1 -> inclusive
    fragment = "\n".join([header] + data_rows)

    df = pd.read_csv(io.StringIO(fragment), sep=",", skipinitialspace=True)
    # First column is the date key (unnamed in source)
    date_col = df.columns[0]
    df = df.rename(columns={date_col: "Date"})
    # Strip whitespace from column names
    df.columns = [c.strip() for c in df.columns]
    return df


def _to_period_index(df: pd.DataFrame) -> pd.DataFrame:
    """Convert the YYYYMM 'Date' column to a monthly PeriodIndex."""
    df = df.copy()
    df["Date"] = df["Date"].astype(str).str.strip()
    # Keep only rows that look like YYYYMM (6-digit)
    mask = df["Date"].str.fullmatch(r"\d{6}")
    df = df[mask].copy()
    df["YYYYMM"] = df["Date"].astype(int)
    df["Period"] = pd.PeriodIndex(
        pd.to_datetime(df["Date"], format="%Y%m"), freq="M"
    )
    return df


def _clean_numeric(df: pd.DataFrame, value_cols: list[str]) -> pd.DataFrame:
    """Coerce value columns to float and replace French missing codes with NaN."""
    df = df.copy()
    for col in value_cols:
        df[col] = pd.to_numeric(df[col], errors="coerce")
    df[value_cols] = df[value_cols].replace(MISSING_CODES, np.nan)
    return df


# --------------------------------------------------------------------------- #
# Step 1: Load, clean, align
# --------------------------------------------------------------------------- #
def load_and_align() -> pd.DataFrame:
    log("=" * 70)
    log("STEP 1: DATA PARSING & ALIGNMENT")
    log("=" * 70)

    # ---- Factors -------------------------------------------------------- #
    factors = _read_block(FACTORS_CSV, FACTORS_HEADER_LINE, FACTORS_DATA_START, FACTORS_DATA_END)
    log(f"[factors] raw parsed shape: {factors.shape}; columns: {list(factors.columns)}")
    factors = _to_period_index(factors)
    factor_value_cols = ["Mkt-RF", "SMB", "HML", "RF"]
    factors = _clean_numeric(factors, factor_value_cols)
    factors_range = (factors["YYYYMM"].min(), factors["YYYYMM"].max())
    log(f"[factors] date range: {factors_range[0]} .. {factors_range[1]}; rows: {len(factors)}")

    # ---- Portfolios (Value-Weighted monthly) ---------------------------- #
    ports = _read_block(PORTFOLIOS_CSV, PORT_HEADER_LINE, PORT_DATA_START, PORT_DATA_END)
    log(f"[portfolios] raw parsed shape: {ports.shape}")
    port_cols = [c for c in ports.columns if c != "Date"]
    log(f"[portfolios] {len(port_cols)} portfolio columns detected")
    ports = _to_period_index(ports)
    ports = _clean_numeric(ports, port_cols)
```

```python
    ports_range = (ports["YYYYMM"].min(), ports["YYYYMM"].max())
    log(f"[portfolios] date range: {ports_range[0]} .. {ports_range[1]}; rows: {len(ports)}")

    # Validate corner labels exist exactly as spelled
    for label in (SMALL_VALUE, BIG_GROWTH):
        if label not in port_cols:
            raise ValueError(
                f"Corner portfolio '{label}' not found. Available: {port_cols}"
            )
    log(f"[portfolios] corner labels verified: Small-Value='{SMALL_VALUE}', Big-Growth='{BIG_GROWTH}'")

    # ---- Merge on Period ------------------------------------------------ #
    factor_keep = factors[["Period", "YYYYMM"] + factor_value_cols]
    port_keep = ports[["Period", SMALL_VALUE, BIG_GROWTH]]
    merged = pd.merge(factor_keep, port_keep, on="Period", how="inner")
    log(f"[merge] inner-join shape (full overlap): {merged.shape}")

    # ---- Filter to sample period ---------------------------------------- #
    sample = merged[(merged["YYYYMM"] >= PERIOD_START) & (merged["YYYYMM"] <= PERIOD_END)].copy()
    sample = sample.sort_values("Period").reset_index(drop=True)
    log(f"[sample] period {PERIOD_START}..{PERIOD_END}: {len(sample)} rows")
    log(f"[sample] first={sample['YYYYMM'].iloc[0]}, last={sample['YYYYMM'].iloc[-1]}")

    # ---- Validation ----------------------------------------------------- #
    # Contiguity: every month present, no gaps
    expected_periods = pd.period_range(
        start=pd.Period("1963-07", freq="M"), end=pd.Period("2023-12", freq="M"), freq="M"
    )
    missing_periods = set(expected_periods) - set(sample["Period"])
    if missing_periods:
        raise ValueError(f"Missing months in sample: {sorted(str(p) for p in missing_periods)}")

    n_missing = int(sample[factor_value_cols + [SMALL_VALUE, BIG_GROWTH]].isna().sum().sum())
    log(f"[validate] total missing values in sample block: {n_missing}")
    if n_missing != 0:
        raise ValueError(f"Sample contains {n_missing} missing values; expected 0.")

    if len(sample) != EXPECTED_T:
        raise ValueError(f"Sample size is {len(sample)}; expected exactly {EXPECTED_T}.")
    log(f"[validate] SAMPLE SIZE VERIFIED: T = {len(sample)} months (expected {EXPECTED_T}) OK")

    # Persist aligned panel for provenance
    out = sample.copy()
    out["Date"] = out["Period"].astype(str)
    out = out[["Date", "YYYYMM", "Mkt-RF", "SMB", "HML", "RF", SMALL_VALUE, BIG_GROWTH]]
    out.to_csv(DATA_DIR / "analysis_data.csv", index=False)
    log(f"[save] aligned panel -> {DATA_DIR / 'analysis_data.csv'}")

    return sample


# ------------------------------------------------------------------------ #
# Step 2: Factor analysis (HML, SMB)
# ------------------------------------------------------------------------ #
def nw_mean_tstat(series: pd.Series, lags: int = NW_LAGS) -> tuple[float, float, float]:
    """Newey-West mean, HAC standard error, and t-statistic via intercept-only OLS.

    Regresses the series on a constant; the intercept equals the sample mean and
    the HAC (Newey-West) standard error gives the autocorrelation/heteroskedasticity
    robust t-statistic of the mean.
    """
    y = series.astype(float).values
    X = np.ones((len(y), 1))
    model = sm.OLS(y, X).fit(cov_type="HAC", cov_kwds={"maxlags": lags})
    mean = float(model.params[0])
    se = float(model.bse[0])
    tstat = float(model.tvalues[0])
    return mean, se, tstat
```

```python


def factor_analysis(sample: pd.DataFrame) -> pd.DataFrame:
    log("=" * 70)
    log("STEP 2: FACTOR ANALYSIS (HML, SMB) - Newey-West, 6 lags")
    log("=" * 70)

    rows = []
    for factor in ["HML", "SMB"]:
        mean_m, se_m, t_m = nw_mean_tstat(sample[factor], NW_LAGS)
        ann_mean = mean_m * 12.0
        pval = 2.0 * (1.0 - _norm_cdf(abs(t_m)))
        rows.append(
            {
                "Factor": factor,
                "Monthly_Mean_pct": mean_m,
                "Annualized_Mean_pct": ann_mean,
                "NW_SE_monthly": se_m,
                "NW_tstat": t_m,
                "p_value": pval,
                "N_months": len(sample),
                "NW_lags": NW_LAGS,
            }
        )
        log(
            f"  {factor}: monthly mean={mean_m:.4f}%  annualized={ann_mean:.4f}%  "
            f"NW_SE={se_m:.4f}  t={t_m:.4f}  p={pval:.4g}"
        )

    df = pd.DataFrame(rows)
    df.to_csv(RESULTS_DIR / "factor_analysis.csv", index=False)
    log(f"[save] factor analysis -> {RESULTS_DIR / 'factor_analysis.csv'}")
    return df


def _norm_cdf(x: float) -> float:
    """Standard normal CDF (t-stats use large-sample normal approximation, T=726)."""
    from math import erf, sqrt

    return 0.5 * (1.0 + erf(x / sqrt(2.0)))


# --------------------------------------------------------------------------- #
# Step 3: Portfolio regressions (CAPM & FF3)
# --------------------------------------------------------------------------- #
def run_regressions(sample: pd.DataFrame) -> tuple[pd.DataFrame, dict]:
    log("=" * 70)
    log("STEP 3: PORTFOLIO REGRESSIONS (CAPM & FF3) - Newey-West, 6 lags")
    log("=" * 70)

    portfolios = {"Small-Value": SMALL_VALUE, "Big-Growth": BIG_GROWTH}

    # Excess returns
    data = sample.copy()
    for name, col in portfolios.items():
        data[f"{name}_excess"] = data[col] - data["RF"]

    mkt = data["Mkt-RF"].values
    smb = data["SMB"].values
    hml = data["HML"].values

    result_rows = []
    fitted_store = {}  # for potential downstream use

    for name, col in portfolios.items():
        y = data[f"{name}_excess"].values

        # --- CAPM ---
```

```python
        X_capm = sm.add_constant(mkt)
        capm = sm.OLS(y, X_capm).fit(cov_type="HAC", cov_kwds={"maxlags": NW_LAGS})
        _record_model(result_rows, name, "CAPM", capm, ["const", "Mkt-RF"])
        log(
            f"  [{name}] CAPM: alpha_m={capm.params[0]:.4f}% (ann={capm.params[0]*12:.4f}%) "
            f"t_alpha={capm.tvalues[0]:.3f}  beta_Mkt={capm.params[1]:.4f}  "
            f"adjR2={capm.rsquared_adj:.4f}"
        )

        # --- FF3 ---
        X_ff3 = sm.add_constant(np.column_stack([mkt, smb, hml]))
        ff3 = sm.OLS(y, X_ff3).fit(cov_type="HAC", cov_kwds={"maxlags": NW_LAGS})
        _record_model(result_rows, name, "FF3", ff3, ["const", "Mkt-RF", "SMB", "HML"])
        log(
            f"  [{name}] FF3 : alpha_m={ff3.params[0]:.4f}% (ann={ff3.params[0]*12:.4f}%) "
            f"t_alpha={ff3.tvalues[0]:.3f}  b_Mkt={ff3.params[1]:.4f}  "
            f"b_SMB={ff3.params[2]:.4f}  b_HML={ff3.params[3]:.4f}  adjR2={ff3.rsquared_adj:.4f}"
        )

    reg_df = pd.DataFrame(result_rows)
    reg_df.to_csv(RESULTS_DIR / "regression_results.csv", index=False)
    log(f"[save] regression results -> {RESULTS_DIR / 'regression_results.csv'}")
    return reg_df, data


def _record_model(rows: list, portfolio: str, model_name: str, fit, term_names: list[str]) -> None:
    """Append one row per estimated coefficient with NW t-stats and model fit."""
    for i, term in enumerate(term_names):
        label = "Alpha" if term == "const" else term
        coef = float(fit.params[i])
        tstat = float(fit.tvalues[i])
        pval = float(fit.pvalues[i])
        annualized = coef * 12.0 if label == "Alpha" else np.nan
        rows.append(
            {
                "Portfolio": portfolio,
                "Model": model_name,
                "Term": label,
                "Coefficient": coef,
                "Annualized_Alpha_pct": annualized,
                "NW_tstat": tstat,
                "p_value": pval,
                "Adj_R2": float(fit.rsquared_adj),
                "N_months": int(fit.nobs),
                "NW_lags": NW_LAGS,
            }
        )


# ------------------------------------------------------------------------- #
# Step 4: Visualizations
# ------------------------------------------------------------------------- #
def make_figures(sample: pd.DataFrame, reg_df: pd.DataFrame) -> None:
    log("=" * 70)
    log("STEP 4: VISUALIZATIONS")
    log("=" * 70)

    plt.rcParams.update(
        {
            "font.family": "sans-serif",
            "font.size": 10,
            "axes.linewidth": 0.8,
            "figure.dpi": 100,
        }
    )

    # ---- Figure 1: Cumulative returns ---------------------------------- #
    # Compound percent-format monthly returns: cumulative product of (1 + r/100).
```

```python
    dates = sample["Period"].dt.to_timestamp()
    series_map = {
        "HML": sample["HML"],
        "SMB": sample["SMB"],
        "Small-Value (SMALL HiBM)": sample[SMALL_VALUE],
        "Big-Growth (BIG LoBM)": sample[BIG_GROWTH],
    }

    fig, ax = plt.subplots(figsize=(11, 6))
    for lbl, ser in series_map.items():
        cum = (1.0 + ser.values / 100.0).cumprod()
        ax.plot(dates, cum, label=lbl, linewidth=1.3)
    ax.set_yscale("log")
    ax.set_title(
        "Cumulative Growth of $1 (Jul 1963 - Dec 2023)\nFama-French Factors & Corner Portfolios",
        fontsize=12,
    )
    ax.set_xlabel("Date")
    ax.set_ylabel("Cumulative value (log scale, $1 initial)")
    ax.legend(loc="upper left", frameon=True, fontsize=9)
    ax.grid(True, which="both", alpha=0.3)
    fig.tight_layout()
    fig.savefig(FIGURES_DIR / "cumulative_returns.png", dpi=300, bbox_inches="tight")
    fig.savefig(FIGURES_DIR / "cumulative_returns.pdf", bbox_inches="tight")
    plt.close(fig)
    log(f"[fig] cumulative returns -> {FIGURES_DIR / 'cumulative_returns.png'}")

    # ---- Figure 2: Factor loadings (betas) ----------------------------- #
    # Compare factor loadings across CAPM and FF3 for both portfolios.
    loadings = reg_df[reg_df["Term"] != "Alpha"].copy()
    portfolios = ["Small-Value", "Big-Growth"]
    factors = ["Mkt-RF", "SMB", "HML"]

    fig, axes = plt.subplots(1, 2, figsize=(13, 5.5), sharey=True)
    for ax, port in zip(axes, portfolios):
        sub = loadings[loadings["Portfolio"] == port]
        capm_vals = [
            _get_loading(sub, "CAPM", f) for f in factors
        ]
        ff3_vals = [
            _get_loading(sub, "FF3", f) for f in factors
        ]
        x = np.arange(len(factors))
        w = 0.38
        b1 = ax.bar(x - w / 2, capm_vals, w, label="CAPM", color="#4C72B0")
        b2 = ax.bar(x + w / 2, ff3_vals, w, label="FF3", color="#C44E52")
        ax.axhline(0, color="black", linewidth=0.6)
        ax.set_xticks(x)
        ax.set_xticklabels(factors)
        ax.set_title(port)
        ax.set_xlabel("Factor")
        ax.grid(True, axis="y", alpha=0.3)
        for bars in (b1, b2):
            for rect in bars:
                h = rect.get_height()
                if not np.isnan(h):
                    ax.annotate(
                        f"{h:.2f}",
                        (rect.get_x() + rect.get_width() / 2, h),
                        ha="center",
                        va="bottom" if h >= 0 else "top",
                        fontsize=8,
                    )
    axes[0].set_ylabel("Factor loading (beta)")
    axes[0].legend(loc="best")
    fig.suptitle(
        "Factor Loadings by Model: CAPM vs Fama-French 3-Factor\n(Newey-West robust estimates, 6 lags)",
        fontsize=12,
```

```
    )
    fig.tight_layout(rect=(0, 0, 1, 0.95))
    fig.savefig(FIGURES_DIR / "factor_loadings.png", dpi=300, bbox_inches="tight")
    fig.savefig(FIGURES_DIR / "factor_loadings.pdf", bbox_inches="tight")
    plt.close(fig)
    log(f"[fig] factor loadings -> {FIGURES_DIR / 'factor_loadings.png'}")


def _get_loading(sub: pd.DataFrame, model: str, factor: str) -> float:
    row = sub[(sub["Model"] == model) & (sub["Term"] == factor)]
    if row.empty:
        return np.nan
    return float(row["Coefficient"].iloc[0])


# ------------------------------------------------------------------------- #
# Main
# ------------------------------------------------------------------------- #
def main() -> int:
    for d in (RESULTS_DIR, FIGURES_DIR, LOGS_DIR, DATA_DIR):
        d.mkdir(parents=True, exist_ok=True)

    log("Fama-French Asset Pricing Analysis")
    log(f"Session: {SESSION_DIR}")
    log(f"Sample period: {PERIOD_START} .. {PERIOD_END} (target T={EXPECTED_T})")
    log("")

    sample = load_and_align()
    log("")
    fa = factor_analysis(sample)
    log("")
    reg_df, _ = run_regressions(sample)
    log("")
    make_figures(sample, reg_df)
    log("")
    log("=" * 70)
    log("ANALYSIS COMPLETE")
    log("=" * 70)
    log(f"Factor analysis rows: {len(fa)}")
    log(f"Regression result rows: {len(reg_df)}")

    # Write log file
    log_path = LOGS_DIR / "analysis_log.txt"
    with open(log_path, "w", encoding="utf-8") as fh:
        fh.write("\n".join(_log_lines) + "\n")
    print(f"[save] log -> {log_path}", flush=True)
    return 0


if __name__ == "__main__":
    sys.exit(main())
```

**Listing 1:** `analyze_assets.py` — full estimation pipeline.

# C Regression output tables (CSV contents)

factor_analysis.csv

```
Factor,Monthly_Mean_pct,Annualized_Mean_pct,NW_SE_monthly,NW_tstat,p_value,N_months,NW_lags
HML,0.29616,3.55388,0.13333,2.22128,0.02633,726,6
SMB,0.17259,2.07107,0.11712,1.47365,0.14057,726,6
```

regression_results.csv

```
Portfolio,Model,Term,Coefficient,Annualized_Alpha_pct,NW_tstat,p_value,Adj_R2,N,lags
Small-Value,CAPM,Alpha, 0.48676, 5.84115, 2.75004,5.96e-03,0.57888,726,6
Small-Value,CAPM,Mkt-RF, 1.06828,        , 23.44838,1.37e-121,0.57888,726,6
Small-Value,FF3, Alpha, 0.16015, 1.92179, 2.51046,1.21e-02,0.90722,726,6
Small-Value,FF3, Mkt-RF, 0.94408,        , 51.21085,0.00e+00,0.90722,726,6
Small-Value,FF3, SMB,    1.10753,        , 22.13785,1.37e-108,0.90722,726,6
Small-Value,FF3, HML,    0.69713,        , 22.92882,2.40e-116,0.90722,726,6
Big-Growth,CAPM, Alpha, 0.02200, 0.26399, 0.30677,7.59e-01,0.88362,726,6
Big-Growth,CAPM, Mkt-RF, 0.98709,        , 51.80670,0.00e+00,0.88362,726,6
Big-Growth,FF3,  Alpha, 0.16908, 2.02895, 3.79509,1.48e-04,0.94627,726,6
Big-Growth,FF3,  Mkt-RF, 0.98562,        , 75.97457,0.00e+00,0.94627,726,6
Big-Growth,FF3,  SMB,   -0.23669,        ,-11.84033,2.41e-32,0.94627,726,6
Big-Growth,FF3,  HML,   -0.35585,        ,-16.90893,3.87e-64,0.94627,726,6
```

## References


[1] Donald W. K. Andrews. Heteroskedasticity and autocorrelation consistent covariance matrix estimation. *Econometrica*, 59(3):817–858, 1991. doi: 10.2307/2938229.

[2] Clifford S. Asness, Tobias J. Moskowitz, and Lasse Heje Pedersen. Value and momentum everywhere. *The Journal of Finance*, 68(3):929–985, 2013. doi: 10.1111/jofi.12021.

[3] Rolf W. Banz. The relationship between return and market value of common stocks. *Journal of Financial Economics*, 9(1):3–18, 1981. doi: 10.1016/0304-405X(81)90018-0.

[4] Sanjoy Basu. Investment performance of common stocks in relation to their price-earnings ratios: A test of the efficient market hypothesis. *The Journal of Finance*, 32(3):663–682, 1977. doi: 10.1111/j.1540-6261.1977.tb01979.x.

[5] Fischer Black, Michael C. Jensen, and Myron Scholes. The capital asset pricing model: Some empirical tests. In Michael C. Jensen, editor, *Studies in the Theory of Capital Markets*, pages 79–121. Praeger, New York, 1972.

[6] Mark M. Carhart. On persistence in mutual fund performance. *The Journal of Finance*, 52(1): 57–82, 1997. doi: 10.1111/j.1540-6261.1997.tb03808.x.

[7] John H. Cochrane. Presidential address: Discount rates. *The Journal of Finance*, 66(4):1047–1108, 2011. doi: 10.1111/j.1540-6261.2011.01671.x.

[8] Eugene F. Fama. Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2):383–417, 1970. doi: 10.2307/2325486.

[9] Eugene F. Fama and Kenneth R. French. The cross-section of expected stock returns. *The Journal of Finance*, 47(2):427–465, 1992. doi: 10.1111/j.1540-6261.1992.tb04398.x.

[10] Eugene F. Fama and Kenneth R. French. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56, 1993. doi: 10.1016/0304-405X(93)90023-5.

[11] Eugene F. Fama and Kenneth R. French. Multifactor explanations of asset pricing anomalies. *The Journal of Finance*, 51(1):55–84, 1996. doi: 10.1111/j.1540-6261.1996.tb05202.x.

[12] Eugene F. Fama and Kenneth R. French. The capital asset pricing model: Theory and evidence. *Journal of Economic Perspectives*, 18(3):25–46, 2004. doi: 10.1257/0895330042162430.

[13] Eugene F. Fama and Kenneth R. French. A five-factor asset pricing model. *Journal of Financial Economics*, 116(1):1–22, 2015. doi: 10.1016/j.jfineco.2014.10.010.

[14] Eugene F. Fama and James D. MacBeth. Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, 81(3):607–636, 1973. doi: 10.1086/260061.

[15] Kenneth R. French. Data library. Tuck School of Business, Dartmouth College, 2024. URL `https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`. Data files created from the CRSP database. Accessed via `https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html`.

[16] Michael R. Gibbons, Stephen A. Ross, and Jay Shanken. A test of the efficiency of a given portfolio. *Econometrica*, 57(5):1121–1152, 1989. doi: 10.2307/1913625.

[17] Campbell R. Harvey, Yan Liu, and Heqing Zhu. ... and the cross-section of expected returns. *The Review of Financial Studies*, 29(1):5–68, 2016. doi: 10.1093/rfs/hhv059.

[18] Kewei Hou, Chen Xue, and Lu Zhang. Digesting anomalies: An investment approach. *The Review of Financial Studies*, 28(3):650–705, 2015. doi: 10.1093/rfs/hhu068.

[19] Michael C. Jensen. The performance of mutual funds in the period 1945–1964. *The Journal of Finance*, 23(2):389–416, 1968. doi: 10.1111/j.1540-6261.1968.tb00815.x.

[20] Josef Lakonishok, Andrei Shleifer, and Robert W. Vishny. Contrarian investment, extrapolation, and risk. *The Journal of Finance*, 49(5):1541–1578, 1994. doi: 10.1111/j.1540-6261.1994.tb04772.x.

[21] John Lintner. The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics*, 47(1):13–37, 1965. doi: 10.2307/1924119.

[22] Andrew W. Lo and A. Craig MacKinlay. Data-snooping biases in tests of financial asset pricing models. *The Review of Financial Studies*, 3(3):431–467, 1990. doi: 10.1093/rfs/3.3.431.

[23] R. David McLean and Jeffrey Pontiff. Does academic research destroy stock return predictability? *The Journal of Finance*, 71(1):5–32, 2016. doi: 10.1111/jofi.12365.

[24] Robert C. Merton. An intertemporal capital asset pricing model. *Econometrica*, 41(5):867–887, 1973. doi: 10.2307/1913811.

[25] Jan Mossin. Equilibrium in a capital asset market. *Econometrica*, 34(4):768–783, 1966. doi: 10.2307/1910098.

[26] Whitney K. Newey and Kenneth D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987. doi: 10.2307/1913610.

[27] Robert Novy-Marx. The other side of value: The gross profitability premium. *Journal of Financial Economics*, 108(1):1–28, 2013. doi: 10.1016/j.jfineco.2013.01.003.

[28] Mitchell A. Petersen. Estimating standard errors in finance panel data sets: Comparing approaches. *The Review of Financial Studies*, 22(1):435–480, 2009. doi: 10.1093/rfs/hhn053.

[29] Barr Rosenberg, Kenneth Reid, and Ronald Lanstein. Persuasive evidence of market inefficiency. *The Journal of Portfolio Management*, 11(3):9–16, 1985. doi: 10.3905/jpm.1985.409007.

[30] William F. Sharpe. Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, 19(3):425–442, 1964. doi: 10.1111/j.1540-6261.1964.tb02865.x.