

Factor Synergy Frontier: A Mixture Design Approach to Weighted-Scoring Stock Selection

Abstract

Multi-factor stock-selection strategies require not only the choice of relevant signals but also the allocation of importance across them. Yet conventional weighting approaches typically produce particular weight vectors without systematically characterizing how portfolio performance varies over the constrained factor-weight space. This study develops a mixture-design framework for factor allocation in weighted-scoring stock selection. Using five signals—price-to-book ratio (PB), return on equity (ROE), market capitalization (MV), Beta, and prior-quarter return scored in the reversal direction (MOM)—the framework combines a five-component simplex-centroid design, historical portfolio backtesting, second-order Scheffé response-surface modeling, and constrained optimization. Thirty-one structured factor-weight combinations are evaluated using quarterly data for CSI 300 stocks, and continuous surfaces are estimated for portfolio return, risk, and the Sharpe ratio.

The empirical results show that portfolio outcomes vary materially with the joint composition of factor weights. The estimated surfaces exhibit economically meaningful nonlinearities across factor combinations, indicating that the consequences of a factor allocation cannot be inferred from standalone factor performance alone. Optimization of the fitted return and risk surfaces further shows that the return-maximizing factor composition changes as the admissible level of portfolio risk varies, generating a Factor Synergy Frontier in factor-weight space. In addition, period-specific Sharpe-ratio optimization reveals substantial shifts in the location of the optimal allocation across subperiods, providing descriptive evidence that the estimated optimal factor composition varies across the realized subperiods.

These findings position factor weighting as a constrained and potentially time-varying allocation problem rather than as a one-time choice of fixed signal weights. The study contributes an integrated framework for mapping and optimizing factor compositions and extends the analysis of time-varying factor performance from individual signals to optimal multi-factor allocations.

Keywords: factor allocation; mixture design; response-surface modeling; Factor Synergy Frontier; factor rotation; multi-factor investing

1. Introduction

1.1 Background and Motivation

Factor investing has become a central approach in empirical asset pricing and quantitative portfolio management. A large body of research shows that stock returns are associated with systematic characteristics related to valuation, profitability, firm size, market risk, and prior-return behavior. As the number of documented predictors has expanded, investment practice has increasingly shifted from reliance on a single signal toward multi-factor stock-selection frameworks.

In a multi-factor strategy, however, selecting the factor set is only part of the portfolio-design problem. When several signals are combined into a weighted-scoring rule, the relative weights assigned to those signals determine how stock characteristics are aggregated into rankings. Different weight vectors can therefore produce different stock selections, portfolio compositions, and performance outcomes even when the underlying factor universe remains unchanged.

This observation makes factor weighting an economically consequential allocation problem rather than a purely technical implementation choice. The weights are jointly constrained: assigning greater importance to one factor necessarily reduces the importance available to others. Moreover, because the composite score determines stock rankings, the performance effect of changing one weight may depend on the complete configuration of the remaining weights. A factor that appears attractive in isolation need not receive a large weight in an optimal multi-factor strategy, while a factor with modest standalone performance may contribute through its role within a broader combination.

The problem also has a temporal dimension. If the relationship between factor composition and portfolio performance changes across market environments, a weighting scheme that performs well in one period may not remain optimal in another. Thus, the relevant question is not simply which factors matter, but how factor importance should be allocated under different investment objectives and whether preferred allocations remain stable through time.

These considerations motivate a framework that treats factor weights as a constrained system and evaluates portfolio outcomes across alternative combinations. The present study develops such a framework by combining mixture design, historical backtesting, response-surface modeling, and constrained optimization. Using five factor signals—PB, ROE, MV, Beta, and MOM—the analysis examines how factor allocations shape portfolio return, risk, and risk-adjusted performance, and whether optimal allocations vary across risk constraints and subperiods.

1.2 Research Gap

Despite extensive research on factor investing and quantitative stock selection, the allocation of weights across multiple factor signals remains insufficiently understood. Existing studies provide substantial evidence on which characteristics are associated with returns, but this evidence does not by itself determine how heterogeneous factors should be combined within a weighted-scoring stock-selection model.

Three related limitations motivate the present study. First, factor weights are jointly constrained and economically interdependent. Increasing the weight assigned to one signal necessarily reduces the weight available to others, while changes in the weight vector may alter stock rankings and portfolio composition. Evaluating factors primarily through standalone performance therefore provides limited information about the consequences of alternative multi-factor allocations.

Second, conventional weighting procedures typically produce a particular weighting scheme—such as equal, judgment-based, or historically estimated weights—without systematically characterizing portfolio performance over the broader feasible weight space. As a result, the relationship between factor composition and portfolio return, risk, or risk-adjusted performance remains only partially observed.

Third, factor performance may vary through time, but evidence of time variation in individual factor returns does not establish whether the optimal combination of factors within a multi-factor ranking model is stable. This leaves open a portfolio-level question: whether the factor-weight allocation associated with superior performance shifts across periods.

These limitations point to the need for a framework that treats factor weights as a constrained allocation system, models portfolio outcomes across alternative weight combinations, and permits optimization across both investment objectives and time periods. The present study develops such a framework by combining mixture design, historical backtesting, response-surface modeling, and constrained optimization.

1.3 Research Questions

Based on the preceding motivation and research gaps, this study addresses three research questions:

RQ1. How does the composition of factor weights affect portfolio return, risk, and risk-adjusted performance within a constrained multi-factor scoring system?

This question examines whether portfolio outcomes depend on the joint configuration of factor weights, including departures from purely linear blending across factor combinations.

RQ2. How do return-maximizing factor allocations change as the admissible level of portfolio risk varies?

This question investigates whether the preferred factor composition is objective-dependent and whether the resulting optimal allocations can be characterized through the **Factor Synergy Frontier**. Factor synergy is an allocation-level concept describing the joint portfolio-performance consequences of combining factor signals within a constrained weighting system. Unlike the classical **efficient frontier**, which optimizes asset weights, the Factor Synergy Frontier operates in factor-weight space and characterizes optimal stock-selection rules under alternative risk constraints.

RQ3. Do Sharpe-maximizing factor allocations remain stable across time, or do they shift across distinct regions of the factor-weight space?

This question evaluates the temporal stability of portfolio-level optimal factor allocations and provides the basis for assessing evidence of factor rotation.

1.4 Research Contributions

This study contributes to the literature in three ways.

Contribution 1: Mixture Design for Factor Allocation

Yeh and Liu (2020) applied mixture design to the search for optimal weights in weighted-scoring stock-selection models. Building on that approach, the present study maps fitted return, volatility, and risk-adjusted performance over the factor-weight space. It then examines how return-maximizing compositions change across volatility constraints and how estimated Sharpe-maximizing compositions differ across realized subperiods. The contribution is this extension of the weighting framework, rather than the first application of mixture design to stock selection.

Contribution 2: Factor Synergy Frontier

The study introduces the concept of the Factor Synergy Frontier. Unlike the traditional efficient frontier, which describes optimal allocations among assets, the proposed frontier describes optimal allocations among factors. By maximizing portfolio return under alternative risk constraints, the frontier reveals how optimal factor compositions vary across risk constraints and provides a factor-allocation representation of the return–risk trade-off.

Contribution 3: Factor Rotation from a Portfolio Perspective

Most existing studies examine factor rotation by analyzing the performance of individual factors. In contrast, this study evaluates factor rotation from the perspective of optimal multi-factor portfolios. By estimating Sharpe-maximizing factor allocations across four subperiods, the analysis reveals substantial changes in optimal factor combinations over time. This portfolio-level perspective provides new evidence regarding the dynamic nature of factor investing and the temporal instability of optimal factor allocations.

2. Literature Review

2.1 Multi-Factor Investing

Factor investing is based on the premise that systematic differences in stock returns can be related to observable firm characteristics, risk exposures, and patterns in prior market behavior. Early asset-pricing research emphasized broad sources of cross-sectional return variation, most notably size and value. Fama and French (1992, 1993) showed that firm size and book-to-market characteristics capture important dimensions of average stock returns beyond the market factor, establishing a foundation for modern multi-factor investing. More recent research has examined whether the prolonged underperformance of conventional value measures reflects a disappearance of the value effect or limitations in how value is measured (Arnott et al., 2021). Haugen and Baker (1996) also examined determinants of expected stock returns and reported common patterns across major equity markets.

Subsequent research expanded the factor universe beyond size and valuation. Profitability emerged as an important dimension when Novy-Marx (2013) documented the predictive relevance of gross profitability and Fama and French (2015) incorporated profitability into their five-factor model. At the same time, evidence on systematic risk challenged the simple prediction that higher market beta should necessarily be associated with superior investment performance. The low-risk literature documents that low-beta or low-volatility securities may deliver unexpectedly favorable risk-adjusted outcomes (Black et al., 1972; Frazzini & Pedersen, 2014).

A separate strand of research emphasizes return patterns associated with investor behavior and short-horizon price dynamics. Lehmann (1990) and Lo and MacKinlay (1990) documented short-term reversal patterns, under which recent relative losers may subsequently outperform recent winners. In contrast, Jegadeesh and Titman (1993) documented profits from buying past winners and selling past losers; the small-preferred prior-quarter return signal used here follows a reversal-oriented ranking direction. These short-term reversal patterns are commonly related to temporary price pressure, market microstructure effects, or investor overreaction. The reversal-oriented prior-return signal therefore represents a conceptually different source of information from valuation, profitability, size, and systematic risk.

Taken together, the literature suggests that multi-factor strategies combine signals drawn from heterogeneous economic dimensions rather than multiple interchangeable measures of the same underlying characteristic. In the present study, PB represents valuation, ROE represents profitability, MV represents firm size, Beta represents systematic risk, and MOM represents prior-return dynamics. The purpose of this factor set is not to provide an exhaustive taxonomy of return predictors, but to construct a heterogeneous multi-factor system in which the consequences of alternative weight allocations can be examined.

This distinction is important because evidence that several factors are individually relevant does not determine how they should be combined. Factors may differ in return potential, risk characteristics, and compatibility with other signals. Moreover, when multiple factors jointly determine stock rankings, the investment outcome depends not only on which factors are included but also on the relative weights assigned to them. The multi-factor literature therefore motivates a separate allocation question: conditional on a chosen factor set, how should importance be distributed across heterogeneous factor signals?

The present study addresses this question using a five-factor weighted-scoring framework. The operational definitions and preferred scoring directions of the factors are specified in Section 3.2, while the systematic allocation of weights is developed through the mixture-design methodology described in Section 3.3.

2.2 Weighted-Scoring Stock Selection

Weighted-scoring models provide a general framework for combining multiple stock characteristics into a single ranking criterion. Let $Z_{ij,t}$ denote the quarterly cross-sectional percentile-rank score of stock j on factor i , transformed according to the factor's prespecified preferred direction, and let w_i denote the corresponding factor weight. The composite score can be written as

$$Score_{j,t}(\mathbf{w}) = \sum_{i=1}^5 w_i Z_{ij,t}. \quad (1)$$

Subject to

$$\sum_{i=1}^k w_i = 1, \quad w_i \geq 0. \quad (2)$$

Stocks are then ranked according to the composite score, and portfolio membership is determined by the resulting ordering. The factor weights therefore affect investment outcomes indirectly through the stock-selection mechanism rather than through direct allocation of capital across securities.

Existing weighting approaches differ in how the vector $\mathbf{w}$ is determined. Equal weighting imposes symmetry across signals, judgment-based methods encode investor or expert preferences, and data-driven procedures estimate weights from historical information. These approaches differ substantially in implementation, but they share a common objective: converting a multidimensional set of stock characteristics into a one-dimensional ranking rule.

The central issue is that the weighting problem is inherently joint. Because the weights are constrained to sum to one, increasing the importance assigned to one factor necessarily reduces the importance available to others. A weight therefore has no

independent meaning outside the complete allocation vector. Moreover, when the composite score determines stock rankings, changes in factor weights can alter portfolio membership nonlinearly. The performance consequences of a given factor weight may therefore depend on the weights assigned to the remaining signals.

This feature distinguishes factor allocation from the simpler task of assessing factor relevance. Evidence that a characteristic predicts returns does not determine the weight it should receive in a multi-factor ranking model. Likewise, a factor that performs well individually need not receive a large optimal weight when combined with competing signals. The relevant empirical object is therefore the performance of the factor-weight vector as a whole.

Although prior research has developed numerous methods for estimating or assigning factor importance, comparatively less attention has been devoted to mapping the performance consequences of alternative weight combinations over the constrained allocation space. In particular, conventional weighting procedures generally produce one weighting vector at a time and therefore provide limited information about how portfolio return, risk, or risk-adjusted performance changes across neighboring or competing allocations.

This limitation motivates the approach adopted in the present study. Rather than treating the factor weights as separately adjustable parameters, the study models them as jointly constrained components of a fixed total. The resulting problem is naturally represented on a simplex, allowing structured factor-weight combinations to be evaluated and their portfolio outcomes modeled as continuous response surfaces. Section 2.3 introduces the mixture-design methodology used for this purpose.

2.3 Mixture Design and Response-Surface Modeling

Mixture-design methodology was developed for problems in which multiple components jointly constitute a fixed total and the response depends on their relative proportions. Unlike conventional experimental designs, mixture designs recognize that component levels are inherently dependent: increasing the proportion of one component necessarily reduces the proportions available to others. This structure has been formalized through simplex-based experimental designs and Scheffé response models (Scheffé, 1958; Cornell, 2011).

A central advantage of mixture design is that it treats the complete composition as the experimental object. Rather than varying each component independently, the methodology evaluates structured combinations over the feasible mixture space. Response-surface modeling can then be used to approximate how an outcome changes across that space, including departures from purely linear blending among components. This makes the approach particularly useful when the objective is not merely to identify

whether individual components matter, but to examine how alternative compositions affect the response.

Mixture methods have been applied extensively in fields such as chemical formulation, food science, pharmaceutical development, and materials engineering, where ingredients are naturally subject to fixed-total constraints (Cornell, 2011; Myers et al., 2016). The underlying statistical structure, however, is not limited to physical mixtures. Any allocation problem in which nonnegative decision weights sum to a fixed total can exhibit the same compositional dependence.

Factor weighting in a weighted-scoring stock-selection model has precisely this structure. Once the factor set is specified, the weights represent relative allocations of importance across signals and are constrained by a common total. A larger weight on one factor necessarily reduces the weight available to competing factors. Moreover, the investment consequences of a weight vector arise from the resulting stock rankings and portfolio composition, so nearby allocations may generate different portfolio outcomes even when the underlying factor universe is unchanged.

This correspondence suggests a natural role for mixture design in factor-allocation research. A structured experimental design can sample informative regions of the feasible weight space, historical backtesting can evaluate the portfolio outcomes associated with those allocations, and response-surface models can approximate the relationship between factor composition and investment performance. The resulting fitted surface can then support analysis of return, risk, and risk-adjusted performance over allocations that were not directly included among the original experimental points.

Despite this methodological fit, mixture-design approaches remain uncommon in factor-investing research. Existing studies typically focus on factor discovery, asset-pricing models, standalone factor performance, or procedures that produce a particular set of estimated weights. Comparatively less attention has been given to treating the constrained factor-weight space itself as an experimental domain and modeling portfolio performance as a response to alternative factor compositions.

The present study applies mixture-design and response-surface methodology to this allocation problem. The contribution is not simply to use an alternative weighting rule, but to construct an empirical mapping from factor-weight compositions to portfolio outcomes and to use that mapping for subsequent optimization. The specific simplex-centroid design and Scheffé models employed in the empirical analysis are defined in Sections 3.3 and 3.5.

2.4 Portfolio Optimization and Time-Varying Factor Allocation

Portfolio optimization traditionally concerns the allocation of capital across assets. Markowitz (1952) formalized the return–risk trade-off by identifying portfolios that maximize expected return for a given level of risk or minimize risk for a given expected return. Subsequent developments have refined the estimation, robustness, and implementation of such asset-allocation problems, but the underlying decision variables generally remain portfolio holdings or asset weights.

Factor-based investment strategies introduce a related but distinct allocation problem. In a weighted-scoring stock-selection model, the decision variables are the relative weights assigned to factor signals. These weights do not allocate capital directly; instead, they determine how multiple stock characteristics are aggregated into rankings and thereby influence portfolio membership. Consequently, optimization in factor-weight space concerns the design of the stock-selection rule rather than the direct composition of asset holdings.

This distinction creates a natural connection between portfolio optimization and factor allocation. If alternative factor-weight vectors generate different return and risk outcomes, the preferred allocation may depend on the investor’s objective and admissible risk level. A factor combination that is attractive under return maximization need not remain optimal under tighter risk constraints or under a risk-adjusted objective. The relevant problem is therefore not simply to identify a single “best” factor, but to characterize how preferred factor compositions vary across investment objectives.

A separate literature documents substantial time variation in factor performance. Research on factor momentum and related dynamic strategies suggests that factor returns may exhibit persistence and that recent factor performance can contain information about subsequent outcomes (Asness et al., 2019; Ehsani & Linnainmaa, 2022). More broadly, this evidence challenges the assumption that the relative attractiveness of factor exposures is constant through time.

However, time variation in individual factor returns does not directly resolve the allocation problem faced by a multi-factor stock-selection model. When several signals jointly determine stock rankings, the optimal weight assigned to one factor may depend on the weights of the others, the resulting portfolio composition, and the performance objective being optimized. Changes in standalone factor performance therefore need not translate one-for-one into changes in optimal multi-factor allocations.

This distinction reveals an important gap between the factor-timing literature and the portfolio-level allocation problem. Existing research provides evidence that factor performance can vary through time, but comparatively less attention has been given to whether the optimal location in a constrained factor-weight space remains stable across periods. The relevant empirical question is therefore not only whether individual factors strengthen or weaken, but whether the factor combination associated with maximum portfolio-level performance shifts across time.

The present study connects these two strands of literature by examining factor allocation in both objective and temporal dimensions. The Factor Synergy Frontier characterizes how return-maximizing factor compositions vary under alternative risk constraints, while the subperiod analysis evaluates the temporal stability of Sharpe-maximizing allocations. The formal optimization procedures are specified in Section 3.6.

2.5 Research Gap and Study Positioning

The preceding literature reveals a gap at the intersection of multi-factor investing, weighted-scoring stock selection, mixture-design methodology, and portfolio optimization. Existing factor research provides extensive evidence on individual return predictors and time-varying factor performance, while weighted-scoring approaches offer practical mechanisms for combining multiple signals. However, these strands do not fully address how portfolio performance varies over the constrained space of factor-weight combinations.

In particular, three issues remain insufficiently integrated. First, evidence on individual factor relevance does not determine how importance should be allocated across heterogeneous signals within a joint ranking rule. Second, conventional weighting procedures generally identify particular weight vectors without explicitly mapping return, risk, and risk-adjusted performance over the broader feasible allocation space. Third, evidence of time-varying factor returns does not establish whether portfolio-level optimal factor combinations remain stable across periods.

This study addresses these issues by treating factor weights as mixture components, using structured backtests to estimate continuous performance surfaces, and optimizing those surfaces across both risk constraints and subperiods. In this way, the study positions factor allocation as a constrained, interactive, and potentially time-varying portfolio-design problem. The specific research questions and contributions arising from this positioning are stated in Sections 1.3 and 1.4.

3. Methodology

3.1 Research Framework

Figure 1 summarizes the empirical framework of the study. The analysis links factor specification, experimental design, portfolio backtesting, response-surface estimation, and optimization within a unified sequence.

The five factor weights are treated as mixture components subject to a constant-sum constraint. A simplex-centroid design is then used to generate a structured set of experimental weight combinations, each of which defines a distinct weighted-scoring stock-selection rule. Historical backtesting provides portfolio-performance

observations for these design points, and second-order Scheffé models are subsequently estimated to transform the discrete backtesting results into continuous response surfaces over the factor-weight space.

The fitted surfaces are then used for two complementary analyses. The first examines how return-maximizing factor allocations change under alternative risk constraints, yielding the Factor Synergy Frontier. The second evaluates whether Sharpe-maximizing factor allocations remain stable across subperiods. The detailed procedures are presented in Sections 3.2–3.6.




Figure 1. The overall research framework.

3.2 Factor Construction

Five stock-level factors are used to construct the weighted-scoring models: PB, ROE, MV, Beta, and MOM. These factors represent valuation, profitability, firm size, systematic risk, and prior-return information, respectively. Their literature foundations are reviewed in Section 2.1; the present section specifies their operational definitions and scoring directions.

Before aggregation, each factor is transformed into a comparable score so that a higher transformed score consistently represents a more preferred stock according to the corresponding signal. PB, Beta, and MOM are treated as small-preferred factors, whereas ROE and MV are treated as large-preferred factors. The resulting factor scores are subsequently combined using the mixture weights defined in Section 3.3.

3.2.1 PB (Small Preferred)

PB is defined as the ratio of market price to book value per share:

$$PB = \frac{\text{Market Price per Share}}{\text{Book Value per Share}}. \quad (3)$$

Lower PB values indicate more favorable valuation under the scoring rule. PB is therefore treated as a small-preferred factor, so stocks with lower PB values receive higher transformed scores.

3.2.2 ROE (Large Preferred)

ROE is defined as net income relative to shareholders' equity:

$$ROE = \frac{\text{Net Income}}{\text{Shareholders' Equity}}. \quad (4)$$

Higher ROE values indicate greater profitability under the scoring rule. ROE is therefore treated as a large-preferred factor, so stocks with higher ROE values receive higher transformed scores.

3.2.3 MV (Large Preferred)

MV denotes market capitalization:

$$MV = \text{Share Price} \times \text{Shares Outstanding}. \quad (5)$$

MV is treated as a large-preferred factor in the present weighted-scoring specification. Accordingly, stocks with larger market capitalization receive higher transformed scores.

3.2.4 Beta (Small Preferred)

Beta measures the sensitivity of a stock's return to market movements. In general form,

$$\text{Beta}_j = \frac{\text{Cov}(R_j, R_m)}{\text{Var}(R_m)}, \quad (6)$$

where R_j denotes the return on stock j and R_m denotes the market return. Beta is treated as a small-preferred factor, so stocks with lower estimated Beta values receive higher transformed scores.

3.2.5 MOM (Short-Term Reversal; Small Preferred)

The MOM factor is based on the stock return observed over the previous quarter. Let $R_{j,t-1}$ denote the prior-quarter return of stock j . The signal is defined operationally by a small-preferred scoring direction: stocks with lower prior-quarter returns receive higher transformed scores.

Thus, MOM is the label retained for the prior-quarter return variable, while its small-preferred scoring direction implements a short-term reversal signal. It should not be interpreted as conventional positive momentum. The label and scoring direction are maintained throughout the analysis.

Table 1 summarizes the five factors, their operational definitions, and their preferred scoring directions.

Table 1. Factor Definitions and Preferred Scoring Directions

Factor Definition		Prespecified Direction	Percentile-Score Mapping	Key References
PB	Price-to-book ratio	Small preferred	Lower raw value → higher score	Fama & French (1992,1993); Lakonishok et al. (1994)
ROE	Net income / shareholders' equity	Large preferred	Higher raw value → higher score	Novy-Marx (2013); Fama & French (2015)
MV	Market capitalization	Large preferred	Higher raw value → higher score	Banz (1981); Fama & French (1993)
Beta	Systematic market-risk measure	Small preferred	Lower raw value → higher score	Sharpe (1964); Frazzini & Pedersen (2014)
MOM	Previous-quarter stock return	Small preferred	Lower raw value → higher score	Lehmann (1990); Lo & MacKinlay (1990)

3.2.6 Factor-Score Transformation

Because the five factors are measured in different units and scales, their raw values are not aggregated directly. Instead, each factor is transformed at every quarterly portfolio-formation date into a cross-sectional percentile-rank score on a common 0–100 scale. The transformation is performed separately for each factor within the contemporaneous investment universe.

Let $X_{ij,t}$ denote the raw value of factor i for stock j at portfolio-formation date t , and let N_t denote the number of eligible stocks in the cross section. For a large-preferred factor, the percentile-rank score is defined so that a larger raw factor value receives a higher score:

$$Z_{ij,t}^{(+)} = 100 \frac{r_{ij,t} - 1}{N_t - 1}, \quad (7)$$

where $r_{ij,t}$ is the ascending cross-sectional rank of $X_{ij,t}$ among the eligible stocks at date t . Under this transformation, the lowest-ranked observation receives a score of 0 and the highest-ranked observation receives a score of 100.

For a small-preferred factor, the scoring direction is reversed:

$$Z_{ij,t}^{(-)} = 100 \frac{N_t - r_{ij,t}}{N_t - 1}. \quad (8)$$

Accordingly, the smallest raw factor value receives a score of 100 and the largest receives a score of 0. In the present study, ROE and MV are large-preferred factors, whereas PB, Beta, and MOM are small-preferred factors.

The transformation is repeated independently at each quarterly portfolio-formation date. Thus, $Z_{ij,t}$ measures the relative cross-sectional standing of stock j on factor i at time t , rather than the absolute magnitude of the raw factor value. This procedure places all five signals on a common 0–100 scale while preserving their prespecified preferred directions.

The composite score of stock j under factor-weight vector $\mathbf{w}$ is then

$$Score_{j,t}(\mathbf{w}) = \sum_{i=1}^5 w_i Z_{ij,t}, \quad (9)$$

subject to

$$\sum_{i=1}^5 w_i = 1, \quad w_i \geq 0. \quad (10)$$

Stocks with higher composite scores are ranked more favorably in the subsequent portfolio-selection procedure.

3.3 Mixture Design

3.3.1 Factor Weights as Mixture Components

The decision variables in the weighted-scoring model are the weights assigned to the five factors: PB, ROE, MV, Beta, and MOM. Let

$$\mathbf{w} = (w_{PB}, w_{ROE}, w_{MV}, w_{Beta}, w_{Reversal})^T \quad (11)$$

denote the factor-weight vector. The weights satisfy

$$\sum_{i=1}^5 w_i = 1, \quad w_i \geq 0, \quad (12)$$

so that the feasible region is the four-dimensional simplex

$$\mathcal{W} = \left\{ \mathbf{w} \in \mathbb{R}^5 : w_i \geq 0, \sum_{i=1}^5 w_i = 1 \right\}. \quad (13)$$

These constraints distinguish the factor-allocation problem from a conventional unconstrained regression-design problem. Increasing the weight assigned to one factor necessarily reduces the weight available to one or more of the remaining factors. The five weights are therefore jointly dependent components of a fixed total and can be treated naturally as mixture components.

Each feasible point $\mathbf{w} \in \mathcal{W}$ defines a distinct weighted-scoring stock-selection rule. The purpose of the experimental design is therefore to select a structured set of factor-weight combinations that spans the simplex and permits the relationship between factor allocation and portfolio performance to be modeled from a manageable number of backtests.

Figure 2 illustrates the geometric concept underlying the simplex-centroid design using a three-component example. The feasible factor-weight region is represented by a simplex because all component weights are nonnegative and constrained to sum to one. The seven design points comprise the three vertices, three edge centroids, and the overall centroid. Although the empirical analysis employs five factors, the same geometric principle extends directly to the five-component simplex, generating the 31 structured design points used in this study.




Figure 2. Geometric Interpretation of the Three-Component Simplex Mixture Design

3.3.2 Simplex-Centroid Design

A five-component simplex-centroid design is employed to generate the experimental factor-weight combinations. In a k -component simplex-centroid design, one design point is included for every nonempty subset of the k components. For a subset containing m factors, the active factors receive equal weights of $1/m$, while all remaining factors receive zero weight.

The total number of design points is therefore

$$\sum_{m=1}^k \binom{k}{m} = 2^k - 1.$$

With $k = 5$, the present design contains

$$2^5 - 1 = 31$$

factor-weight combinations. These consist of:

$$\binom{5}{1} = 5$$

single-factor vertices,

$$\binom{5}{2} = 10$$

two-factor edge centroids,

$$\binom{5}{3} = 10$$

three-factor subset centroids,

$$\binom{5}{4} = 5$$

four-factor subset centroids, and

$$\binom{5}{5} = 1$$

overall centroid.

Thus,

$$5 + 10 + 10 + 5 + 1 = 31.$$

For example, a two-factor design point involving PB and ROE is represented by

$$(0.5, 0.5, 0, 0, 0),$$

whereas a three-factor design point involving PB, ROE, and Reversal is represented by

$$\left(\frac{1}{3}, \frac{1}{3}, 0, 0, \frac{1}{3}\right).$$

The overall centroid assigns equal weight to all five factors:

$$(0.2, 0.2, 0.2, 0.2, 0.2).$$

Each of the 31 design points defines one experimental weighted-scoring model and is evaluated using the common backtesting procedure described in Section 3.4. The resulting portfolio-performance observations are subsequently used to estimate the continuous response surfaces in Section 3.5.

3.4 Backtesting Procedure

The backtesting procedure evaluates the 31 factor-weight combinations generated by the simplex-centroid design. Portfolios are formed and rebalanced quarterly. At each quarterly portfolio-formation date t , the eligible investment universe is defined using the historical CSI 300 constituent membership applicable at that date. Thus, only stocks that were actual constituents of the CSI 300 at the corresponding formation date are included in the cross-sectional ranking and portfolio-selection procedure.

For each eligible stock, the five factor signals—PB, ROE, MV, Beta, and MOM (short-term reversal)—are constructed using information available under the temporal-alignment rules specified below. At each quarterly formation date, the factor values are transformed separately into cross-sectional percentile-rank scores according to the prespecified scoring directions defined in Section 3.2.6. ROE and MV are large-preferred factors, whereas PB, Beta, and MOM are small-preferred factors. The resulting scores are expressed on a common 0–100 scale, with higher scores consistently representing more preferred stocks.

For a given factor-weight vector

$$\mathbf{w} = (w_{\text{PB}}, w_{\text{ROE}}, w_{\text{MV}}, w_{\text{Beta}}, w_{\text{MOM}})^T \quad (14)$$

the composite score of stock j at portfolio-formation date t is calculated as

$$Score_{j,t}(\mathbf{w}) = \sum_{i=1}^5 w_i Z_{ij,t}, \quad (15)$$

subject to

$$\sum_{i=1}^5 w_i = 1, \quad w_i \geq 0. \quad (16)$$

Each of the 31 simplex-centroid design points therefore defines a distinct weighted-scoring stock-selection rule. At each quarterly formation date, eligible stocks are ranked in descending order according to $\text{Score}_{j,t}(\mathbf{w})$. The top 10% of stocks are selected, assigned equal portfolio weights, and held until the next quarterly rebalancing date. This procedure is repeated independently for all 31 factor-weight combinations, producing a separate backtested portfolio-return series for each design point.

To prevent look-ahead bias, all factor inputs are aligned with the portfolio-formation date using information available under the study's timing conventions. For accounting-based variables, a four-quarter reporting lag is imposed. Specifically, accounting information associated with quarter $t-4$ is used in forming portfolios at quarter t . Accordingly, the accounting components used in the construction of PB and ROE at portfolio-formation quarter t are based on information lagged by four quarters. This conservative timing rule prevents contemporaneous or subsequently reported accounting information from entering the stock-ranking procedure.

The MOM signal is defined as the stock return over the immediately preceding quarter. Thus, at portfolio-formation quarter t , the MOM input is based on the return realized in quarter $t-1$. Because MOM is scored in the small-preferred direction, stocks with lower prior-quarter returns receive higher percentile-rank scores. Beta is estimated using return information available before the corresponding portfolio-formation date, thereby excluding subsequent return observations from factor construction. MV is measured using market information available at the portfolio-formation date.

The use of historical CSI 300 constituent membership preserves the contemporaneous investment opportunity set throughout the backtest. At each formation date, the eligible universe is reconstructed from the constituent list applicable in that quarter rather than from a fixed ex post constituent list. This procedure avoids incorporating future index membership into historical portfolio formation and reduces survivorship bias arising from the retrospective use of later constituent information.

Following portfolio formation, portfolio returns are measured over the subsequent quarterly holding period until the next rebalancing date. For each of the 31 design-point portfolios, the resulting return series is used to calculate three performance measures: mean portfolio return, the standard deviation of portfolio returns, and the Sharpe ratio. These performance measures serve as the response variables in the second-order Scheffé models described in Section 3.5.

For each strategy, mean return is the arithmetic average of its quarterly portfolio returns, and return standard deviation is calculated from the same quarterly return series. Both mean return and return standard deviation are reported in quarterly percentage points. The risk-free return is set to zero in the present calculation. Accordingly, the

reported Sharpe ratio equals the mean quarterly portfolio return divided by the standard deviation of quarterly portfolio returns.

3.5 Performance Modeling

The portfolio outcomes obtained from the 31 simplex-centroid design points are used to estimate continuous response surfaces over the feasible factor-weight space. Separate second-order Scheffé mixture models are fitted for mean return, return standard deviation, and the Sharpe ratio.

Let

$$\mathbf{w} = (w_{\text{PB}}, w_{\text{ROE}}, w_{\text{MV}}, w_{\text{Beta}}, w_{\text{Reversal}})^{\top} \quad (17)$$

denote the factor-weight vector, where

$$\sum_{i=1}^5 w_i = 1, \quad w_i \geq 0. \quad (18)$$

For a generic portfolio-performance response Y , the second-order Scheffé model is specified as

$$Y(\mathbf{w}) = \sum_{i=1}^5 \beta_i w_i + \sum_{i<j} \beta_{ij} w_i w_j + \varepsilon, \quad (19)$$

where β_i denotes a linear mixture coefficient, β_{ij} denotes a pairwise interaction coefficient, and ε is the error term.

Accordingly, three separate response models are estimated:

$$R(\mathbf{w}) = \sum_{i=1}^5 \beta_i^{(R)} w_i + \sum_{i<j} \beta_{ij}^{(R)} w_i w_j + \varepsilon_R, \quad (20)$$

$$\sigma(\mathbf{w}) = \sum_{i=1}^5 \beta_i^{(\sigma)} w_i + \sum_{i<j} \beta_{ij}^{(\sigma)} w_i w_j + \varepsilon_{\sigma}, \quad (21)$$

$$S(\mathbf{w}) = \sum_{i=1}^5 \beta_i^{(S)} w_i + \sum_{i<j} \beta_{ij}^{(S)} w_i w_j + \varepsilon_S, \quad (22)$$

where R , σ , and S denote portfolio mean return, return standard deviation, and the Sharpe ratio, respectively.

The fitted models yield the continuous surfaces

$$\hat{R}(\mathbf{w}), \quad \hat{\sigma}(\mathbf{w}), \quad \hat{S}(\mathbf{w}), \quad (23)$$

which approximate portfolio performance throughout the feasible simplex rather than only at the 31 observed design points.

Because the factor weights satisfy a constant-sum constraint, the coefficients should be interpreted within the geometry of the mixture model. The linear mixture

coefficients are not independent marginal effects in the conventional regression sense. Similarly, the pairwise interaction terms capture departures from the corresponding linear blending pattern within the fitted surface. A positive interaction coefficient may indicate that a particular two-factor region lies above the linear blending benchmark, but it should not automatically be interpreted as proof of causal or globally optimal economic synergy.

The complete fitted surfaces, rather than individual coefficients in isolation, provide the basis for the subsequent optimization analyses. The return and risk surfaces are jointly used to construct the Factor Synergy Frontier, while the Sharpe-ratio surface is used to identify full-sample and period-specific risk-adjusted optima.

3.6 Optimization Procedures

3.6.1 Factor Synergy Frontier

For a given admissible volatility level C , the return-maximizing factor allocation is defined as

$$\mathbf{w}^*(C) = \arg \max_{\mathbf{w} \in \mathcal{W}} \widehat{R}(\mathbf{w}) \quad (24)$$

subject to

$$\widehat{\sigma}(\mathbf{w}) \leq C. \quad (25)$$

By solving this problem over a sequence of volatility limits, the analysis obtains a set of optimal factor-weight vectors

$$\{\mathbf{w}^*(C)\}, \quad (26)$$

together with their corresponding predicted return-risk pairs

$$\left(\widehat{\sigma}(\mathbf{w}^*(C)), \widehat{R}(\mathbf{w}^*(C)) \right). \quad (27)$$

The resulting locus is termed the **Factor Synergy Frontier**. Unlike the classical Markowitz efficient frontier, whose decision variables are asset weights, the present frontier is defined over factor weights used in the stock-selection rule. Each point therefore represents a model-implied optimal factor composition under a specified volatility constraint rather than a direct allocation of capital across securities.

In this study, factor synergy is defined at the allocation level rather than by the sign of any individual interaction coefficient. It refers to the joint performance consequences of combining heterogeneous factor signals within a constrained weighting system, where the preferred composition depends on the complete configuration of factor weights and the investor's admissible risk level. Accordingly, the Factor Synergy Frontier is the locus of return-maximizing factor-weight vectors obtained under

alternative volatility constraints. The term does not imply that every pair of factors exhibits positive statistical interaction or causal complementarity.

The frontier is determined by the complete fitted return and risk surfaces under the mixture constraints. Pairwise interaction terms may contribute to the curvature of those surfaces, but no individual coefficient is sufficient to determine the location or composition of the frontier.

3.6.2 Full-Sample Sharpe-Ratio Optimization

To identify a benchmark allocation for risk-adjusted performance, the fitted full-sample Sharpe-ratio surface is optimized over the feasible simplex:

$$\mathbf{w}_{\text{FS}}^* = \arg \max_{\mathbf{w} \in \mathcal{W}} \hat{S}(\mathbf{w}) \quad (28)$$

This optimization differs from the Factor Synergy Frontier problem. The frontier maximizes predicted return conditional on a sequence of risk limits, whereas the present problem directly maximizes the fitted Sharpe ratio. The resulting full-sample optimum therefore provides a distinct benchmark for comparison with the period-specific optima examined below.

3.6.3 Temporal Stability of Optimal Factor Allocations

To evaluate whether the Sharpe-maximizing factor allocation remains stable through time, the sample is divided into four subperiods. For each period t , a separate second-order Scheffé model is estimated for the Sharpe ratio, yielding a period-specific fitted surface

$$\hat{S}_t(\mathbf{w}). \quad (29)$$

The optimal factor allocation for period t is defined as

$$\mathbf{w}_t^* = \arg \max_{\mathbf{w} \in \mathcal{W}} \hat{S}_t(\mathbf{w}). \quad (30)$$

Temporal stability is assessed by comparing the locations of

$$\mathbf{w}_1^*, \mathbf{w}_2^*, \mathbf{w}_3^*, \mathbf{w}_4^* \quad (31)$$

within the common factor-weight simplex. If the optimal vectors remain concentrated in a similar region, the evidence is consistent with a relatively stable risk-adjusted allocation. If they move materially across distinct regions of the simplex, the estimates indicate descriptive variation in the location of the subperiod optima; this comparison alone does not establish a statistically significant change in the underlying population optimum.

In this study, such movement is interpreted as evidence of factor rotation at the level of multi-factor allocations. This interpretation is explicitly *ex post*. The procedure compares optima estimated from realized subperiod data and does not constitute a

forecasting rule for future factor weights. Establishing a predictive factor-timing strategy would require a separate design based on real-time information, rolling or expanding estimation windows, and out-of-sample evaluation.

4. Factor Synergy Frontier

4.1 Estimated Performance Surfaces

The 31 backtesting observations are used to estimate separate second-order Scheffé models for mean return, return standard deviation, and the Sharpe ratio. Table 2 reports the estimated linear mixture coefficients and pairwise interaction coefficients for the three response surfaces.

Across the 31 design points, the fitted mean-return, volatility, and Sharpe-ratio surfaces have R^2 values of 0.945, 0.946, and 0.922, respectively. These statistics describe interpolation fit within the evaluated design points and do not constitute out-of-sample validation.

The coefficient patterns differ substantially across the three performance measures, indicating that factor composition is associated differently with return, risk, and risk-adjusted performance. Because the factor weights satisfy a constant-sum constraint, the linear mixture coefficients should not be interpreted as independent marginal effects. Likewise, an interaction coefficient represents a departure from the corresponding linear blending pattern within the fitted mixture surface rather than a standalone measure of economic synergy.

Several interaction coefficients have small conventional OLS p values within the fitted mixture models. In particular, MOM enters positive estimated interactions with each of the other four signals in at least one performance model. The ROE×MOM and PB×MOM terms are prominent in the fitted surfaces. These coefficient patterns describe departures from linear blending among the 31 evaluated design points; they do not, by themselves, establish robust or causal economic synergy.

The three fitted models therefore provide distinct continuous mappings,

$$\widehat{R}(\mathbf{w}), \quad \widehat{\sigma}(\mathbf{w}), \quad \widehat{S}(\mathbf{w}), \quad (32)$$

over the feasible factor-weight simplex. The subsequent optimization analysis uses the complete fitted surfaces rather than individual coefficients in isolation.

Figure 3 provides a conceptual illustration of the proposed Factor Synergy Frontier. Under linear blending, two-factor mixtures would lie on the line segment connecting the corresponding pure-factor outcomes, while mixtures of all five factors would lie within the convex hull of the pure-factor outcomes. However, interaction effects among factors may bend the frontier outward, allowing diversified factor combinations to outperform the linear blending benchmark. The intermediate portfolios therefore

represent potential factor synergies rather than simple weighted averages of individual factor performance. This conceptual relationship provides the motivation for estimating second-order Scheffé mixture models, through which the nonlinear return, risk, and Sharpe-ratio surfaces are examined in the following empirical analysis.

Table 2. Estimated Second-Order Scheffé Models for Portfolio Performance

Term	Mean Return Coef.	t	p	Return STD Coef.	t	p	Sharpe Ratio Coef.	t	p
PB	7.12	10.09	<0.001	30.26	10.02	<0.001	0.24	30.99	<0.001
ROE	10.66	15.11	<0.001	47.41	15.69	<0.001	0.22	29.52	<0.001
MV	5.01	7.11	<0.001	27.22	9.01	<0.001	0.18	23.99	<0.001
Beta	8.18	11.60	<0.001	41.27	13.66	<0.001	0.20	26.12	<0.001
MOM	10.26	14.55	<0.001	57.59	19.07	<0.001	0.18	23.28	<0.001
PB × ROE	2.87	1.01	0.330	5.25	0.43	0.672	0.04	1.33	0.204
PB × MV	-0.64	-0.22	0.830	3.46	0.28	0.780	-0.03	-0.99	0.339
PB × Beta	-5.80	-2.04	0.060	-30.20	-2.48	0.025	0.01	0.38	0.712
PB × MOM	-2.33	-0.82	0.430	-50.57	-4.15	0.001	0.20	6.49	<0.001
ROE × MV	-5.71	-2.01	0.060	-16.76	-1.37	0.188	-0.03	-1.01	0.329
ROE × Beta	9.22	3.24	0.010	49.16	4.03	0.001	-0.02	-0.81	0.431
ROE × MOM	15.30	5.38	<0.001	34.62	2.84	0.012	0.14	4.58	<0.001
MV × Beta	-3.06	-1.08	0.300	-13.35	-1.10	0.290	-0.01	-0.17	0.863
MV × MOM	1.90	0.67	0.510	-9.52	-0.78	0.446	0.10	3.24	0.005
Beta × MOM	4.99	1.76	0.100	2.32	0.19	0.851	0.10	3.20	0.006
R ²		0.945			0.946			0.922	

Note: The signs in parentheses indicate the prespecified scoring directions used in the quarterly percentile-rank transformation, rather than the signs of the estimated coefficients. A positive sign indicates that larger raw factor values receive higher scores, whereas a negative sign indicates that smaller raw factor values receive higher scores. Accordingly, MOM(-) denotes the previous-quarter return signal with a small-preferred scoring direction and therefore represents a reversal-oriented implementation. The reported t statistics and p values are conventional regression quantities and do not adjust for common historical quarters or overlapping portfolio holdings across design points. R² and RMSE measure in-sample fit across the 31 simplex-centroid design points.




Figure 3. Conceptual Illustration of the Factor Synergy Frontier under Non-linear Mixture Blending.

Note: The curves depict the departure from linear blending paths due to pairwise factor interaction effects, illustrating how multi-factor allocations (e.g., $(0.25, 0.5, 0.25)$) can optimize the risk-return profile beyond individual single-factor vertices.

4.2 Construction of the Factor Synergy Frontier

The fitted return and risk surfaces are jointly used to construct the Factor Synergy Frontier. For each admissible volatility level C , the optimization problem identifies the feasible factor-weight vector that maximizes predicted portfolio return subject to the corresponding risk constraint.

Table 3 reports the resulting optimal allocations, while Figure 4 plots the fitted frontier together with the 31 observed backtesting outcomes. The frontier therefore represents a sequence of model-implied optimal factor compositions rather than a set of directly observed portfolios.

At relatively loose volatility limits, the risk constraint is nonbinding and the same return-maximizing solution is obtained. In particular, the 80% and 100% volatility limits yield identical allocations because the predicted volatility of the unconstrained return-maximizing solution is approximately 61.0%. The solution assigns 51.3% to ROE and 48.7% to MOM and produces a predicted quarterly return of approximately 14.28%. As the admissible volatility level is reduced, the constraint becomes binding and the return-maximizing allocation changes. The resulting sequence traces the trade-off between predicted return and predicted risk across the factor-weight space.

Table 3. Optimal Factor Weights Along the Factor Synergy Frontier

Volatility Constraint (%)	Optimal weights					Predicted values (%)	
	PB	ROE	MV	Beta	MOM	Mean Return	Return STD
100.0	0.000	0.513	0.000	0.000	0.487	14.284	61.017
80.0	0.000	0.513	0.000	0.000	0.487	14.284	61.017
50.0	0.197	0.476	0.000	0.000	0.328	12.332	50.000
45.0	0.312	0.377	0.000	0.000	0.311	11.334	45.000
40.0	0.428	0.267	0.000	0.000	0.305	10.291	40.000
35.0	0.537	0.153	0.000	0.000	0.310	9.203	35.000
30.0	0.635	0.041	0.000	0.000	0.324	8.076	30.000
29.0	0.654	0.019	0.000	0.000	0.327	7.846	29.000
28.0	0.681	0.000	0.000	0.000	0.319	7.614	28.000
27.5	0.735	0.000	0.000	0.146	0.120	6.910	27.500
27.2	0.682	0.000	0.000	0.318	0.000	6.200	27.213

Note: Volatility Constraint, Mean Return, and Return STD are all expressed in percentage points at the quarterly frequency. Optimal factor weights are proportions that sum to one.




Figure 4. The backtesting results of the 31 factor weight combinations and the Factor Synergy Frontier.

4.3 Evolution of Optimal Factor Weights

Figure 5 shows how the return-maximizing factor allocation changes along the Factor Synergy Frontier. The movement is systematic but nonlinear, with different factors entering or leaving the optimum as the admissible volatility level changes.

At the high-return end of the frontier, the optimal allocation is concentrated in ROE and MOM. The return-maximizing solution assigns approximately 51.3% to ROE and 48.7% to MOM, with a predicted quarterly return of 14.28% and predicted volatility of approximately 61.0%.

As the volatility limit is tightened from 50% to 30%, PB progressively enters the optimal allocation and its weight rises from 19.7% to 63.5%. Over the same range, the ROE weight declines from 47.6% to 4.1%, while MOM remains within a comparatively narrower range of approximately 30%–33%. Thus, the intermediate region of the frontier is characterized primarily by substitution between PB and ROE while MOM remains an important component.

Near the low-risk end of the frontier, the composition changes again. ROE falls to zero, MOM declines sharply, and Beta enters the optimal solution. At the reported minimum-risk endpoint, the allocation is concentrated in PB and Beta, with weights of 68.2% and 31.8%, respectively.

Overall, the frontier displays a structured transition from a profitability–reversal allocation at the high-return end, through intermediate allocations in which PB increasingly substitutes for ROE, toward a value–low-beta allocation near the minimum-risk region. These results show that the return-maximizing factor composition is conditional on the admissible level of portfolio risk.




Figure 5. The optimal weights along the Factor Synergy Frontier.

4.4 Economic Interpretation of Frontier Composition

The frontier reveals distinct factor-weight patterns across the estimated return–risk trade-off. ROE is concentrated near the high-return end and declines as the volatility constraint tightens, whereas PB becomes increasingly prominent toward the lower-risk region. MOM remains present across much of the frontier before declining near the

minimum-risk endpoint, while Beta enters only near the low-risk boundary. MV receives zero weight in the reported frontier solutions.

These patterns describe the locations of the five signals within the estimated factor-weight surface. They should not be interpreted as universal causal properties of the underlying factors. In particular, the presence of a factor in a low-risk optimum does not establish that the factor independently reduces risk, and a zero optimal weight does not imply that the factor lacks predictive value.

Overall, the frontier indicates that the preferred factor composition changes with the admissible risk level. A broader discussion of objective-dependent factor roles and their implications for multi-factor portfolio design is provided in Section 6.2.

4.5 Full-Sample Sharpe-Ratio Optimization

The Factor Synergy Frontier characterizes return-maximizing factor allocations under alternative volatility constraints. A complementary question is which factor composition maximizes predicted risk-adjusted performance over the full sample. To provide a benchmark for the subsequent temporal analysis, this section optimizes the fitted full-sample Sharpe-ratio surface over the feasible factor-weight simplex.

The optimization problem is

$$\max_{\mathbf{w}} \hat{S}(\mathbf{w}) \quad (33)$$

subject to

$$\sum_{i=1}^5 w_i = 1, \quad w_i \geq 0. \quad (34)$$

The resulting full-sample optimum assigns 64.7% of the total weight to PB and 35.3% to MOM, while ROE, MV, and Beta receive zero weight. The predicted Sharpe ratio is 0.260, with a predicted mean return of 7.69% and a predicted return standard deviation of 28.3%.

This solution differs substantially from the full-sample return-maximizing allocation at the high-return end of the Factor Synergy Frontier, which is concentrated in ROE and MOM. The contrast shows that the preferred factor composition depends on the optimization objective. When predicted return is maximized without a binding low-volatility requirement, the solution is concentrated in profitability and MOM signals; when the fitted Sharpe-ratio surface is optimized directly, PB replaces ROE in the optimal allocation.

The positive weight assigned to MOM under both objectives indicates that MOM remains part of the optimum under distinct full-sample criteria. By contrast, the roles of ROE and PB are objective-dependent: ROE appears in the return-maximizing solution, whereas PB dominates the Sharpe-maximizing solution. These differences

should be interpreted as properties of the complete fitted surfaces under the mixture constraints rather than as consequences of individual coefficients in isolation.

The zero weights assigned to ROE, MV, and Beta in the Sharpe-maximizing solution do not imply that these factors lack predictive or economic relevance. They indicate only that these factors receive zero weight at the numerical optimum of the estimated full-sample Sharpe-ratio surface over the specified feasible simplex.

This full-sample solution serves as a benchmark rather than evidence of a universally stable allocation. Section 5 therefore examines whether the Sharpe-maximizing region remains similar across subperiods or shifts materially through time.

5. Temporal Stability of Sharpe-Maximizing Factor Allocations

5.1 Sharpe-Ratio Models Across Periods

Following the temporal-stability procedure described in Section 3.6.3, separate second-order Scheffé models are estimated for the Sharpe ratio in each of the four subperiods. The estimated coefficients are reported in Table 4.

The coefficient structures vary substantially across periods, showing substantial cross-period variation in the estimated shape of the fitted Sharpe-ratio surface. This variation appears in both the linear mixture coefficients and the pairwise interaction terms. Because the factor weights are subject to the constant-sum constraint, individual coefficients should be interpreted as components of the complete fitted surface rather than as independent marginal effects.

In Period 1, MOM has the largest linear mixture coefficient (0.339), followed by ROE (0.304), PB (0.275), and Beta (0.263). Several interactions involving MOM are also positive, including PB×MOM (0.215), MV×MOM (0.119), and Beta×MOM (0.197). The estimated surface therefore assigns an important role to MOM-related regions of the factor-weight space during this period.

In Period 2, PB, ROE, and MOM have similarly large linear mixture coefficients of 0.271, 0.261, and 0.256, respectively. The MOM coefficient is lower than in Period 1, while ROE×MOM becomes positive (0.113). Most other pairwise interaction coefficients are comparatively small in magnitude. Relative to Period 1, the fitted surface therefore exhibits a different structure, with less widespread positive interactions involving MOM.

Period 3 displays another distinct configuration. ROE has the largest linear mixture coefficient (0.331), followed by MV (0.265) and MOM (0.242). Several interactions involving PB are negative, including PB×ROE (-0.244), PB×MV (-0.326), PB×Beta (-

0.371), and PB×MOM (-0.174). These estimates indicate a substantial reshaping of the fitted surface relative to the first two periods.

The largest coefficient shifts occur in Period 4. PB has the largest positive linear mixture coefficient (0.403), whereas the ROE coefficient becomes negative (-0.228). Several interaction coefficients are large and positive, including PB×ROE (1.021), PB×MV (0.714), PB×Beta (1.400), PB×MOM (0.989), MV×Beta (0.675), and Beta×MOM (0.847). At the same time, ROE×Beta (-0.432) and ROE×MV (-0.217) are negative. Taken together, these estimates indicate pronounced curvature and a major reconfiguration of the fitted Sharpe-ratio surface in the final subperiod.

Overall, Table 4 shows that the relationship between factor composition and predicted risk-adjusted performance changes materially across periods. However, the relative magnitude of an individual coefficient does not by itself determine the optimal factor allocation. Under the mixture constraint, the optimum depends on the complete set of linear and interaction terms evaluated jointly over the feasible simplex. The period-specific optima are therefore examined directly in Section 5.2.

Table 4. Estimated Sharpe-Ratio Models Across Four Subperiods

Term	Period 1	Period 2	Period 3	Period 4
PB	0.275	0.271	0.188	0.403
ROE	0.304	0.261	0.331	-0.228
MV	0.188	0.177	0.265	0.172
BETA	0.263	0.236	0.153	0.252
MOM	0.339	0.256	0.242	0.112
PB × ROE	0.240	-0.038	-0.244	1.021
PB × MV	-0.020	0.020	-0.326	0.714
PB × BETA	0.123	-0.006	-0.371	1.400
PB × MOM	0.215	-0.005	-0.174	0.989
ROE × MV	0.014	-0.084	0.170	-0.217
ROE × BETA	0.024	-0.091	0.057	-0.432
ROE × MOM	-0.009	0.113	0.140	0.047
MV × BETA	0.021	-0.041	-0.088	0.675
MV × MOM	0.119	0.077	-0.082	0.325
BETA × MOM	0.197	0.026	0.047	0.847

5.2 Optimal Sharpe-Ratio Allocations Across Periods

The period-specific Sharpe-ratio surfaces are optimized according to the procedure defined in Section 3.6.3. For each subperiod, the objective is to identify the feasible

factor-weight vector that maximizes the fitted Sharpe ratio over the simplex. The resulting optimal allocations are reported in Table 5 and illustrated in Figure 6.

The period-specific optima are determined by the complete fitted Sharpe-ratio surfaces under the mixture constraints. Although some allocations are broadly consistent with prominent coefficient patterns in Table 4, the optimal weights should not be attributed to individual coefficients in isolation.

In Period 1, the Sharpe-maximizing allocation assigns 61.5% to MOM, 28.7% to PB, and 9.8% to Beta, with a predicted Sharpe ratio of 0.366. The solution is therefore concentrated in a three-factor region of the simplex dominated by MOM.

In Period 2, the optimum shifts to a two-factor allocation consisting of 52.2% ROE and 47.8% MOM. The predicted Sharpe ratio is 0.287. Relative to Period 1, PB and Beta disappear from the optimum, while profitability becomes a major component of the Sharpe-maximizing allocation.

In Period 3, the optimal allocation moves to another region of the simplex, assigning 69.5% to ROE and 30.5% to MV. The predicted Sharpe ratio rises to 0.347. MOM receives zero weight, and the optimum is concentrated entirely in a profitability–size combination.

In Period 4, the optimal allocation changes again. PB receives 55.0% of the total weight, Beta receives 44.1%, and MOM receives a negligible 0.9%. The predicted Sharpe ratio reaches 0.681. Thus, the final-period optimum is concentrated primarily in a value–low-beta region of the factor-weight space.

Taken together, the four solutions occupy markedly different regions of the simplex. The dominant composition changes from a MOM-centered allocation in Period 1, to a ROE–MOM allocation in Period 2, to a ROE–MV allocation in Period 3, and finally to an allocation concentrated in PB and Beta in Period 4. These shifts provide the direct optimization evidence used in Section 5.3 to assess the temporal stability of optimal factor allocations.

Table 5. Optimal Sharpe-Ratio Allocations Across Four Subperiods

Period	PB	ROE	MV	Beta	MOM	Predicted Sharpe Ratio
Period 1	0.287	0.000	0.000	0.098	0.615	0.366
Period 2	0.000	0.522	0.000	0.000	0.478	0.287
Period 3	0.000	0.695	0.305	0.000	0.000	0.347
Period 4	0.550	0.000	0.000	0.441	0.009	0.681

5.3 Evidence of Portfolio-Level Factor Rotation

The period-specific optimization results provide descriptive evidence of substantial variation in the estimated location of the Sharpe-maximizing allocation. As shown in Table 5 and Figure 6, the optimal weight vectors occupy markedly different regions of

the factor-weight simplex across the four subperiods rather than remaining concentrated around a stable long-run composition.

In this study, factor rotation is defined at the portfolio-allocation level as a material shift in the location of the Sharpe-maximizing factor-weight vector across subperiods. This definition differs from approaches that identify rotation through changes in standalone factor returns or through a real-time switching rule among individual factors. Because the five signals jointly determine stock rankings, the relevant object is the optimal weight vector as a whole.

The observed shifts are economically meaningful because they involve changes in the composition of the optimal multi-factor allocation rather than merely fluctuations in individual-factor performance. Signals that receive substantial weight in one period may receive little or no weight in another, while previously absent factors may become central components of the optimum. The evidence therefore indicates that the location of the portfolio-level optimum is not stable across the realized subperiods.

Under the definition adopted here, these movements provide evidence of **portfolio-level factor rotation**. However, the result should be interpreted as **ex post temporal instability** rather than as proof of a predictive factor-timing strategy. The analysis estimates separate Sharpe-ratio surfaces for realized subperiods and compares their corresponding optima; it does not test whether future shifts in the optimal allocation can be predicted using information available at the portfolio-formation date.




Figure 6. The optimal weights of maximizing Sharpe Index in the four periods.

Accordingly, the results show that a single fixed factor-weight vector does not characterize the Sharpe-maximizing allocation uniformly across the sample. Whether the observed temporal variation can be converted into an implementable dynamic

allocation rule remains a separate question requiring explicit forecasting, rolling or expanding estimation windows, and out-of-sample evaluation.

6. Discussion

6.1 Factor Effects Are Conditional Rather Than Isolated

The results indicate that factor importance is conditional on the broader allocation in which a factor is embedded. In a weighted-scoring model, increasing the weight assigned to one signal necessarily reduces the relative importance of competing signals and may alter stock rankings, portfolio membership, and subsequent performance. Accordingly, the economic role of a factor cannot be inferred from its standalone performance alone.

This point is distinct from the claim that every pair of factors exhibits statistical or economic synergy. The estimated response surfaces contain departures from purely linear blending, but individual interaction coefficients should be interpreted within the complete mixture model. A positive pairwise coefficient does not by itself establish causal complementarity, nor does it imply that the corresponding factor pair forms a globally optimal allocation.

The broader implication is that factor evaluation and factor allocation are related but different problems. Evidence that a characteristic predicts returns may justify its inclusion in a candidate factor set, but it does not determine the weight that the characteristic should receive when combined with other signals. For allocation purposes, the relevant object is therefore the factor-weight vector as a whole.

This perspective complements the conventional emphasis on factor discovery. As the number of documented return predictors has expanded, an increasingly important question is how heterogeneous signals should be combined after they have been selected. The mixture-design framework addresses this question by treating the constrained factor-weight space itself as the domain of analysis.

6.2 Factor Roles Depend on the Performance Objective

The empirical results also show that factor roles are objective-dependent. The compositions associated with higher predicted return are not identical to those associated with lower predicted risk or higher predicted risk-adjusted performance. A factor may therefore receive substantial weight under one optimization criterion and little or no weight under another.

This finding cautions against assigning fixed labels such as “return factor,” “defensive factor,” or “risk-reducing factor” without qualification. Within the present sample, some factors are more prominent in particular regions of the fitted performance surfaces, but these patterns are conditional on the model, scoring directions, factor set,

and optimization objective. They should not be interpreted as universal causal properties of the underlying characteristics.

The distinction is especially important in multi-factor stock selection because factor weights influence performance indirectly through the ranking mechanism. Changing the weight vector can alter which stocks enter the portfolio, so the effect of a factor allocation is mediated by changes in portfolio composition. Consequently, the same signal may play different roles when combined with different competing signals or evaluated under different objectives.

The practical implication is that factor allocation should begin with an explicit performance criterion. A weighting scheme designed to maximize predicted return need not be appropriate for an investor prioritizing volatility control or the Sharpe ratio. The empirical question is therefore not whether one factor is universally superior, but which joint composition is preferred under the stated objective and constraints.

6.3 The Factor Synergy Frontier as an Allocation-Space Concept

The Factor Synergy Frontier extends the return–risk allocation problem from asset-weight space to factor-weight space. In the classical portfolio setting, the decision variables are direct holdings in securities. In the present framework, the decision variables are weights assigned to signals within a stock-selection rule. The frontier therefore characterizes optimal ranking systems under alternative volatility constraints rather than direct capital allocations across assets.

This distinction is conceptually important. Factor weights first determine how firm characteristics are aggregated into composite scores; the scores then determine stock rankings and portfolio membership; and the resulting holdings generate return and risk. The mapping from factor weights to portfolio outcomes is therefore mediated by the stock-selection process and need not resemble a simple linear allocation among standalone factor portfolios.

The empirical frontier shows that the return-maximizing factor composition changes as the admissible risk level varies. Its significance lies in this movement of the optimum across the constrained factor-weight space. A single fixed weighting vector does not represent the return-maximizing solution uniformly across different risk tolerances.

The term *synergy* is used here in an allocation-level sense. It refers to the joint portfolio-performance consequences of combining heterogeneous signals within a constrained weighting system, where the preferred composition depends on the complete factor-weight configuration and the investor’s admissible risk level. It does

not imply that every factor pair exhibits positive statistical interaction, causal complementarity, or superior standalone performance.

Accordingly, the frontier is determined by the complete fitted return and risk surfaces under the mixture constraints. Pairwise interaction terms may contribute to surface curvature, but no individual coefficient is sufficient to determine the location of an optimal allocation. This interpretation preserves the conceptual role of the Factor Synergy Frontier while distinguishing it from a coefficient-based definition of synergy.

6.4 Temporal Instability and Portfolio-Level Factor Rotation

The subperiod analysis adds a temporal dimension to the allocation problem. The Sharpe-maximizing weight vectors occupy materially different regions of the factor-weight simplex across the realized subperiods, indicating that the full-sample optimum does not represent a uniformly stable risk-adjusted allocation throughout the sample.

This result extends the discussion of time-varying factor performance from individual signals to multi-factor combinations. Conventional factor-timing research often examines whether the returns of individual factors vary predictably through time. The present analysis addresses a different question: whether the location of the portfolio-level optimum changes when several signals jointly determine stock rankings.

Under the definition adopted in this study, such material movement constitutes portfolio-level factor rotation. The concept refers to shifts in the Sharpe-maximizing factor-weight vector across subperiods rather than to a real-time switching rule among individual factors. This distinction is important because changes in standalone factor performance need not translate one-for-one into changes in the optimal multi-factor composition.

At the same time, the evidence remains *ex post*. The period-specific surfaces are estimated from realized observations within each subperiod, and the corresponding optima are compared after estimation. The analysis documents differences among the *ex post* estimated subperiod optima, without establishing that the underlying optimal allocations differ statistically or that future movements can be forecast.

The implication is consequently more limited, but also more precise, than a claim of successful factor timing. The results question the assumption that one fixed factor-weight vector is uniformly optimal through time and motivate further research on whether changes in the optimal region can be anticipated. Establishing an implementable dynamic strategy would require rolling or expanding estimation, explicit forecasting variables, transaction-cost treatment, and genuine out-of-sample evaluation.

6.5 Practical and Research Implications

The findings have several implications for multi-factor portfolio design. First, factor weighting should be treated as a joint allocation problem rather than as a collection of independent parameter choices. Because the weights are constrained and jointly determine stock rankings, the consequences of a weighting decision depend on the complete composition.

Second, optimization should be aligned with the investor's objective. The factor composition preferred under return maximization can differ from that preferred under direct Sharpe-ratio maximization or tighter volatility constraints. The choice of objective function is therefore an economically substantive part of strategy design rather than a purely technical detail.

Third, the observed temporal instability cautions against assuming that a full-sample optimum is necessarily stable. At the same time, the present evidence does not justify mechanical periodic re-optimization or a predictive timing strategy. Any dynamic implementation would require evidence that future changes in the optimal allocation can be forecast and exploited after accounting for estimation error, turnover, and transaction costs.

From a research perspective, the framework shifts attention from factor discovery alone toward the geometry of the factor-weight space. This opens a broader set of questions concerning the stability, predictability, and economic interpretation of optimal multi-factor compositions.

7. Conclusion

This study examines factor weighting as a constrained allocation problem within a multi-factor stock-selection system. Using five heterogeneous signals—PB, ROE, MV, Beta, and MOM—the analysis integrates mixture design, historical backtesting, response-surface modeling, and constrained optimization to examine how alternative factor compositions are associated with portfolio return, risk, and risk-adjusted performance.

Factor Allocation, Objective Dependence, and Temporal Instability

The results yield three main conclusions. First, portfolio outcomes vary materially across the constrained factor-weight space. The estimated response surfaces contain meaningful departures from purely linear blending, indicating that the consequences of a factor allocation depend on the broader composition in which each signal is embedded. This supports a joint-allocation perspective without implying that every positive interaction coefficient represents causal or universal economic synergy.

Second, optimal factor composition is objective-dependent. The return-maximizing allocation changes as the admissible volatility level varies, producing a

Factor Synergy Frontier in factor-weight space. Unlike the classical efficient frontier, which optimizes direct asset holdings, this frontier characterizes optimal factor weights within a stock-selection rule. The full-sample Sharpe-ratio optimum further differs from the return-maximizing solution, reinforcing the conclusion that preferred factor weights depend on the stated performance objective.

Third, the Sharpe-maximizing region of factor-weight space changes materially across subperiods. Under the definition adopted in this study, these shifts provide descriptive evidence of portfolio-level factor rotation. The finding extends the analysis of time-varying factor performance from individual signals to optimal multi-factor combinations. At the same time, the evidence is *ex post* and should not be interpreted as proof of a predictive factor-timing rule.

Limitations

Several limitations qualify these conclusions. The empirical findings are conditional on the selected equity universe, sample period, five-factor set, scoring directions, portfolio-construction rule, and rebalancing design. The second-order Scheffé specification provides an interpretable approximation to the performance surface but may not capture more complex nonlinear structures. In addition, the subperiod analysis identifies temporal instability without establishing the economic states responsible for the shifts or demonstrating that future movements can be forecast.

Directions for Future Research

Future research can extend the framework across alternative markets, factor universes, scoring transformations, and portfolio rules. More flexible response-surface specifications may also be compared with the second-order mixture model. Most importantly, a prospective design based on rolling or expanding estimation windows, information available at each decision date, transaction costs, and genuine out-of-sample evaluation would allow researchers to test whether changes in the optimal factor-weight region are predictable and economically exploitable.

Concluding Perspective

Overall, the study shows that multi-factor portfolio design involves more than selecting a set of individually relevant signals. Once factors are combined within a common ranking system, their relative weights form a constrained allocation space whose preferred regions depend on the performance objective and may shift through time.

List of Symbols

Symbol	Definition
i	Index of factors, $i=1,\dots,5$
j	Index of stocks
t	Index of quarterly portfolio-formation dates or subperiods, as specified by context
k	Number of factors in the mixture system; $k=5$ in this study
N_t	Number of eligible stocks in the investment universe at formation date t
$X_{ij,t}$	Raw value of factor i for stock j at date t
$r_{ij,t}$	Cross-sectional rank of $X_{ij,t}$ among eligible stocks at date t
$Z_{ij,t}$	Percentile-rank score of stock j on factor i at date t , transformed to the common 0–100 scale according to the prespecified preferred direction
$Z_{ij,t}(+)$	Percentile-rank score for a large-preferred factor
$Z_{ij,t}(-)$	Percentile-rank score for a small-preferred factor
w_i	Weight assigned to factor i
w	Factor-weight vector
W	Feasible factor-weight simplex
$\text{Score}_{j,t}(w)$	Composite weighted score of stock j at date t under factor-weight vector w
R_j	Return on stock j
R_m	Market return
$R_{j,t-1}$	Previous-quarter return of stock j , used as the raw MOM signal
$Y(w)$	Generic portfolio-performance response associated with factor-weight vector w
$R(w)$	Portfolio mean-return response
$\sigma(w)$	Portfolio return-standard-deviation response
$S(w)$	Portfolio Sharpe-ratio response
β_i	Linear mixture coefficient for factor i
β_{ij}	Pairwise interaction coefficient for factors i and j
ε	Error term in the generic response-surface model
C	Admissible portfolio-volatility constraint
$w^*(C)$	Return-maximizing factor-weight vector under volatility constraint C
w_{FS}^*	Full-sample Sharpe-maximizing factor-weight vector
w_t^*	Sharpe-maximizing factor-weight vector for subperiod t

Symbol	Definition
PB	Price-to-book ratio
ROE	Return on equity
MV	Market capitalization
Beta	Systematic market-risk measure
Reversal	Previous-quarter stock return used with a small-preferred scoring direction

References

- Arnott, R. D., Harvey, C. R., Kalesnik, V., & Linnainmaa, J. T. (2021). Reports of value's death may be greatly exaggerated. *Financial Analysts Journal*, 77(1), 44–67.
- Asness, C. S., Frazzini, A., Israel, R., Moskowitz, T. J., & Pedersen, L. H. (2019). Factor momentum everywhere. *Journal of Portfolio Management*, 45(3), 13–36.
- Banz, R. W. (1981). The relationship between return and market value of common stocks. *Journal of Financial Economics*, 9(1), 3–18.
- Black, F., Jensen, M. C., & Scholes, M. (1972). The capital asset pricing model: Some empirical tests. In M. C. Jensen (Ed.), *Studies in the theory of capital markets* (pp. 79–121). Praeger.
- Cornell, J. A. (2011). *Experiments with mixtures: Designs, models, and the analysis of mixture data* (3rd ed.). Wiley.
- Ehsani, S., & Linnainmaa, J. T. (2022). Factor momentum and the momentum factor. *Journal of Finance*, 77(3), 1877–1919.
- Fama, E. F., & French, K. R. (1992). The cross-section of expected stock returns. *Journal of Finance*, 47(2), 427–465.
- Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1), 3–56.
- Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, 116(1), 1–22.
- Frazzini, A., & Pedersen, L. H. (2014). Betting against beta. *Journal of Financial Economics*, 111(1), 1–25.
- Haugen, R. A., & Baker, N. L. (1996). Commonality in the determinants of expected stock returns. *Journal of Financial Economics*, 41(3), 401–439.
- Jegadeesh, N., & Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. *Journal of Finance*, 48(1), 65–91.
- Lakonishok, J., Shleifer, A., & Vishny, R. W. (1994). Contrarian investment, extrapolation, and risk. *Journal of Finance*, 49(5), 1541–1578.

- Lehmann, B. N. (1990). Fads, martingales, and market efficiency. *Quarterly Journal of Economics*, 105(1), 1–28.
- Lo, A. W., & MacKinlay, A. C. (1990). When are contrarian profits due to stock market overreaction? *Review of Financial Studies*, 3(2), 175–205.
- Markowitz, H. (1952). Portfolio selection. *Journal of Finance*, 7(1), 77–91.
- Myers, R. H., Montgomery, D. C., & Anderson-Cook, C. M. (2016). *Response surface methodology: Process and product optimization using designed experiments* (4th ed.). Wiley.
- Novy-Marx, R. (2013). The other side of value: The gross profitability premium. *Journal of Financial Economics*, 108(1), 1–28.
- Scheffé, H. (1958). Experiments with mixtures. *Journal of the Royal Statistical Society, Series B*, 20(2), 344–360.
- Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *Journal of Finance*, 19(3), 425–442.
- Yeh, I. C., & Liu, Y. C. (2020). Discovering optimal weights in weighted-scoring stock-picking models: A mixture design approach. *Financial Innovation*, 6(1), 41.