Portfolio Theory and Applications
Final Project

Factor-Based Long-Short Allocation Models with Beta Constraints

FE630

May 16, 2025

Prepared by:
Srinath Sundar
Vineel Pentrala
Pavan Prakash

Submitted to:
Prof. Papa Momar Ndiaye

Abstract

This report presents the development and comparative analysis of two factor-based long-short allocation models with differing beta constraints. The first strategy employs a conservative approach with beta constrained to $[-0.5, 0.5]$, while the second strategy adopts a more aggressive stance with beta constrained to $[-2, 2]$. We utilize the Fama-French 3-factor model to estimate returns and construct covariance matrices, implementing weekly portfolio rebalancing over the period from March 2007 to December 2024. Our analysis explores the sensitivity of these strategies to different estimation windows (short-term: 40 days, mid-term: 90 days, long-term: 180 days) and risk aversion parameters (0.1, 0.5, 1.0) across various market regimes. We find that estimation window selection significantly impacts strategy performance, with different windows proving optimal during crisis and non-crisis periods. The results provide valuable insights into the behavior of factor-based strategies under different market conditions and offer practical recommendations for portfolio managers regarding the appropriate choice of estimators.

Contents

1 Introduction **3**

2 Theoretical Framework **3**
2.1 Factor Models and Portfolio Optimization 3
2.2 Investment Strategies 4
2.2.1 Strategy I: Conservative Approach 4
2.2.2 Strategy II: Aggressive Approach 4
2.3 Term Structure of Estimators 5

3 Methodology **5**
3.1 Data Description 5
3.2 Market Regimes 6
3.3 Implementation Approach 6
3.3.1 Factor Model Estimation 6
3.3.2 Rolling Window Estimation 6
3.3.3 Portfolio Optimization 6
3.3.4 Performance Evaluation 7

4 Empirical Results **7**
4.1 Overall Performance Comparison 7
4.2 Performance Across Market Regimes 9
4.3 Impact of Risk Aversion Parameter 11
4.4 Drawdown Analysis 12

5 Discussion and Recommendations **16**
5.1 Optimal Estimator Selection 16
5.2 Strategy Selection and Risk Management 17
5.3 Practical Implementation Considerations 18

6 Conclusion **18**

1 Introduction

Portfolio optimization remains a cornerstone of modern financial theory, with practitioners continually seeking robust strategies that can weather various market conditions. In this project, we explore factor-based long-short allocation models with specific beta constraints, aiming to understand how the choice of estimators impacts strategy performance across different market environments.

The primary objective of this study is to build and compare two distinct portfolio optimization strategies: a conservative approach with tight beta constraints (Strategy I) and a more aggressive approach with wider beta constraints (Strategy II). We implement these strategies using the Fama-French 3-factor model, which decomposes returns into market (momentum), size, and value factors. By analyzing the sensitivity of these strategies to the term structure of estimators—short-term (40 days), mid-term (90 days), and long-term (180 days)—we seek to determine the optimal approach for different market conditions.

Our analysis spans from March 2007 to December 2024, encompassing significant market events including the 2008 financial crisis and the COVID-19 pandemic. This extended timeframe allows us to evaluate strategy performance across five distinct regimes: pre-crisis, 2008 crisis, normal market conditions, COVID period, and post-COVID recovery.

Through this comprehensive analysis, we aim to address several key questions:

- How does the choice of estimation window length affect strategy performance?
- Which combination of risk and return estimators performs best during crisis versus non-crisis periods?
- How does the risk aversion parameter influence portfolio characteristics and performance?
- What practical recommendations can be derived for portfolio managers regarding estimator selection?

The remainder of this report is organized as follows: Section 2 presents the theoretical framework underlying our strategies; Section 3 describes our methodology, including data sources and estimation techniques; Section 4 presents our empirical results; Section 5 discusses our findings and offers recommendations; and Section 6 concludes.

2 Theoretical Framework

2.1 Factor Models and Portfolio Optimization

Factor models have become a cornerstone of modern portfolio theory, providing a parsimonious way to capture the systematic risks affecting asset returns. In this project, we employ the Fama-French 3-factor model, which extends the Capital Asset Pricing Model (CAPM) by adding size and value factors to the market factor.

Under the Fama-French 3-factor model, the return of a security is given by:

$$r_i = r_f + \beta_i^3(r_M - r_f) + b_i^s r_{SMB} + b_i^v r_{HML} + \alpha_i + \varepsilon_i \tag{1}$$

where r_f is the risk-free rate, r_M is the market return, r_{SMB} is the return on the size factor (Small Minus Big), r_{HML} is the return on the value factor (High Minus Low), α_i is the security-specific excess return, and ε_i is the residual term. The coefficients β_i^3, b_i^s, and b_i^v represent the security's exposures to the market, size, and value factors, respectively.

In terms of expected returns, this translates to:

$$\rho_i = r_f + \beta_i^3(\rho_M - r_f) + b_i^s \rho_{SMB} + b_i^v \rho_{HML} + \alpha_i \tag{2}$$

It's important to note that the market beta in the Fama-French model (β_i^3) may differ from the traditional CAPM beta (β_i^m), defined as:

$$\beta_i^m = \frac{\text{cov}(r_i, r_M)}{\sigma^2(r_M)} \tag{3}$$

2.2 Investment Strategies

We implement two distinct portfolio optimization strategies with different risk-return trade-offs and beta constraints.

2.2.1 Strategy I: Conservative Approach

Strategy I employs a mean-variance optimization with tight beta constraints:

$$\max_{\omega \in \mathbb{R}^n} \rho^T \omega - \lambda \sqrt{\omega^T \Sigma \omega} \tag{4}$$

$$\text{s.t.} \quad -0.5 \leq \sum_{i=1}^{n} \beta_i^m \omega_i \leq 0.5 \tag{5}$$

$$\sum_{i=1}^{n} \omega_i = 1 \tag{6}$$

$$-2 \leq \omega_i \leq 2 \tag{7}$$

This strategy aims to achieve a portfolio that is relatively market-neutral, with a beta constrained to the interval $[-0.5, 0.5]$. The objective function balances expected return ($\rho^T \omega$) against portfolio volatility ($\sqrt{\omega^T \Sigma \omega}$), with the parameter λ controlling the risk aversion.

2.2.2 Strategy II: Aggressive Approach

Strategy II incorporates a tracking error component and allows for wider beta swings:

$$\max_{\omega \in \mathbb{R}^n} \rho^T \omega - r_{SPY}^T TEV(\omega) - \lambda \sqrt{\omega^T \Sigma \omega} \tag{8}$$

$$\text{s.t.} \quad -2 \leq \sum_{i=1}^{n} \beta_i^m \omega_i \leq 2 \tag{9}$$

$$\sum_{i=1}^{n} \omega_i = 1 \tag{10}$$

$$-2 \leq \omega_i \leq 2 \tag{11}$$

where $TEV(\omega) = \sigma(r_P(\omega) - r_{SPY})$ is the Tracking Error Volatility relative to the S&P 500.

This strategy allows for more significant market exposure, with beta constrained to the wider interval $[-2, 2]$. The inclusion of the tracking error term encourages deviations from the benchmark only when they are expected to yield high returns.

2.3 Term Structure of Estimators

A critical aspect of our analysis is the examination of how the choice of estimation window affects strategy performance. We consider three different window lengths:

- Short-Term (ST): 40 days - captures recent market trends but may be more susceptible to noise
- Mid-Term (MT): 90 days - balances recency with statistical robustness
- Long-Term (LT): 180 days - provides more stable estimates but may be slower to adapt to changing market conditions

These windows are applied to the estimation of both expected returns and covariance matrices, creating a "term structure" of estimators. We hypothesize that different combinations of these estimators will perform optimally under different market conditions.

3 Methodology

3.1 Data Description

Our investment universe consists of 12 ETFs representing diverse asset classes and geographical regions:

1. CurrencyShares Euro Trust (FXE)
2. iShares MSCI Japan Index (EWJ)
3. SPDR GOLD Trust (GLD)
4. Powershares NASDAQ-100 Trust (QQQ)
5. SPDR S&P 500 (SPY)
6. iShares Lehman Short Treasury Bond (SHV)
7. PowerShares DB Agriculture Fund (DBA)
8. United States Oil Fund LP (USO)
9. SPDR S&P Biotech (XBI)
10. iShares S&P Latin America 40 Index (ILF)

11. iShares MSCI Pacific ex-Japan Index Fund (EPP)
12. SPDR DJ Euro Stoxx 50 (FEZ)

We obtain daily price data for these ETFs from Yahoo Finance, covering the period from March 1, 2007, to December 31, 2024. For the Fama-French factors (Market, SMB, HML), we utilize data from Kenneth French's data library.

3.2 Market Regimes

To analyze strategy performance across different market environments, we divide our sample period into five distinct regimes:

- Pre-Crisis: March 1, 2007, to August 31, 2007
- Crisis08: September 1, 2007, to March 31, 2009
- Normal: April 1, 2009, to December 31, 2019
- COVID: January 1, 2020, to June 30, 2021
- Post-COVID: July 1, 2021, to December 31, 2024

These divisions allow us to compare strategy performance during periods of market stress versus relative stability.

3.3 Implementation Approach

Our implementation follows a systematic approach:

3.3.1 Factor Model Estimation

For each ETF, we estimate the Fama-French factor loadings (β_i^3, b_i^s, b_i^v) using linear regression. We also calculate the traditional CAPM beta (β_i^m) for use in the portfolio constraints.

3.3.2 Rolling Window Estimation

To ensure our backtesting is non-anticipative (using only information available at the time of decision-making), we implement a rolling window approach for estimating expected returns and covariance matrices. For each rebalancing date, we use only historical data from the preceding window (40, 90, or 180 days).

3.3.3 Portfolio Optimization

We solve the optimization problems for Strategies I and II on a weekly basis. For Strategy I (convex optimization), we employ the CVXPY solver. For Strategy II (nonlinear optimization), we use the SLSQP method from SciPy.

3.3.4 Performance Evaluation

We evaluate strategy performance using a comprehensive set of metrics:

- Cumulative Return
- Annualized Mean Return and Volatility
- Sharpe Ratio
- Maximum Drawdown
- Value-at-Risk (VaR) and Conditional VaR (CVaR)
- Skewness and Kurtosis

We compute these metrics for the overall period as well as for each market regime, allowing for a nuanced analysis of strategy behavior.

4 Empirical Results

4.1 Overall Performance Comparison

We begin by examining the overall performance of our strategies across different estimation windows and risk aversion parameters.

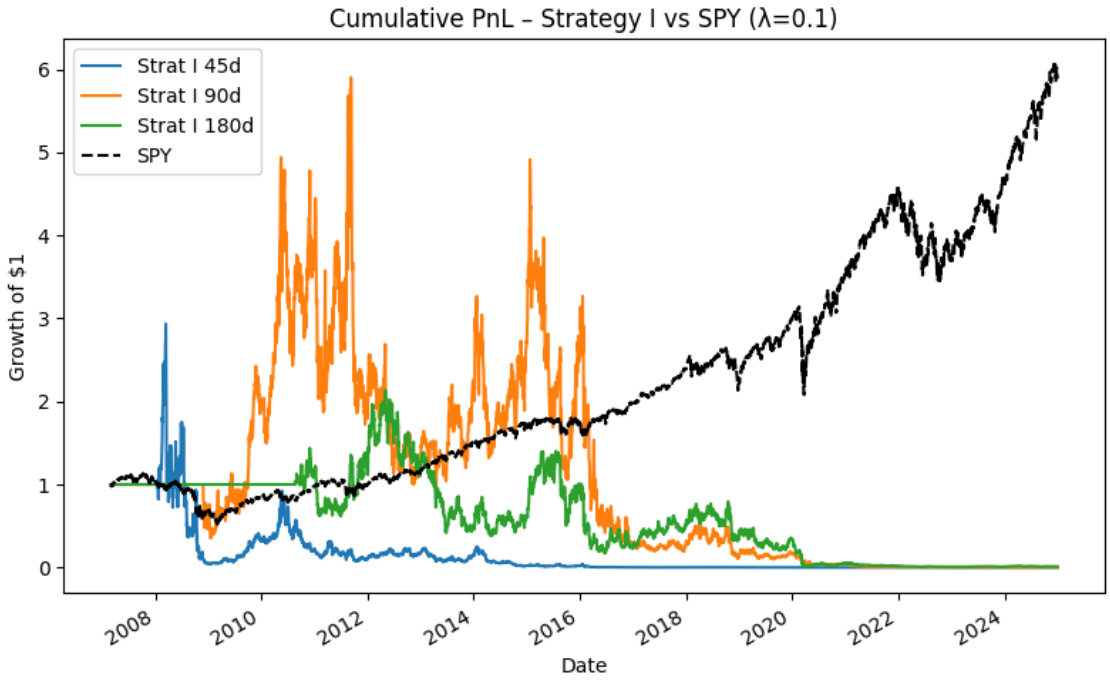


Figure 1: Cumulative Performance of Strategy I vs. SPY ($\lambda = 0.1$)

Figure 1 presents the cumulative returns of Strategy I using different estimation windows (45, 90, and 180 days), compared against the S&P 500 benchmark with a risk aversion parameter of $\lambda = 0.1$. Several interesting patterns emerge from this visualization. The 90-day estimation window (orange line) shows remarkable outperformance during specific periods, particularly in 2012-2014, with cumulative returns nearly 6 times the initial investment. However, all three variants of Strategy I experience substantial declines post-2016, with performance significantly deteriorating by 2020.

It's particularly noteworthy that the 45-day estimation window (blue line) demonstrates strong performance during the early part of the 2008 financial crisis but fails to sustain this advantage over the longer term. The 180-day window (green line) shows more moderate performance peaks but also exhibits substantial drawdowns, particularly evident around 2014-2016.

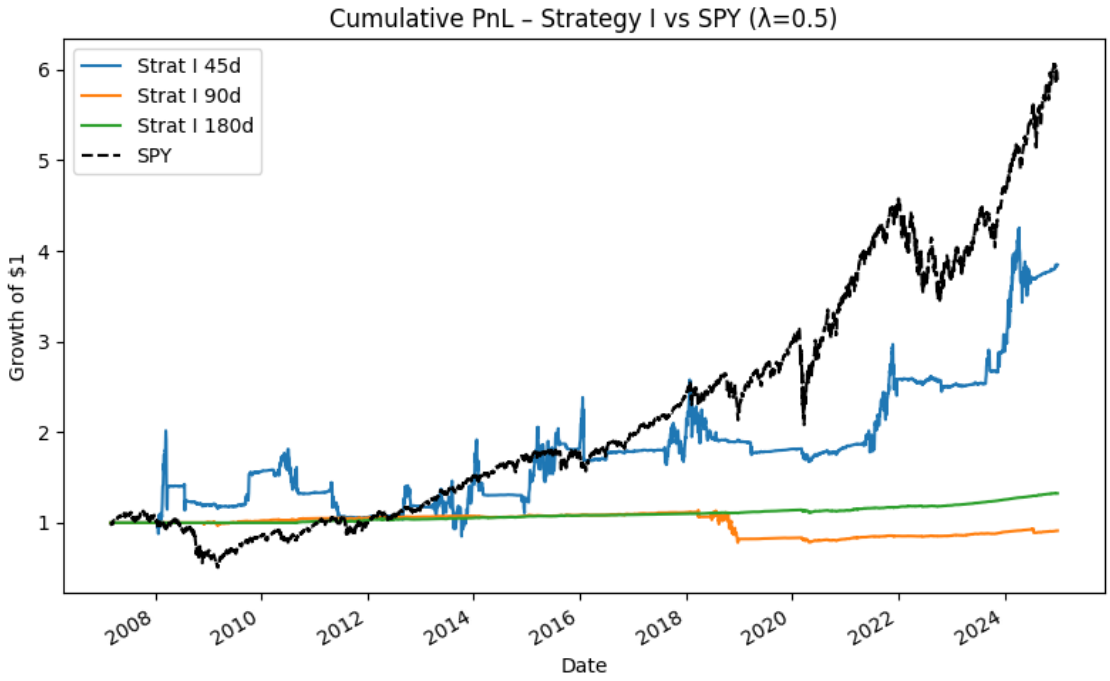


Figure 2: Cumulative Performance of Strategy I vs. SPY ($\lambda = 0.5$)

Figure 2 shows how increasing the risk aversion parameter to $\lambda = 0.5$ impacts Strategy I's performance. With this higher risk aversion, we observe significantly more stable performance, particularly for the 45-day window (blue line), which now shows impressive long-term growth, especially after 2020. This contrasts sharply with its performance at $\lambda = 0.1$, suggesting that the shorter-term estimation window benefits from stronger risk controls.

The 90-day and 180-day windows show more muted performance with the higher risk aversion parameter, indicating that these longer-term estimators become overly conservative when combined with higher λ values.

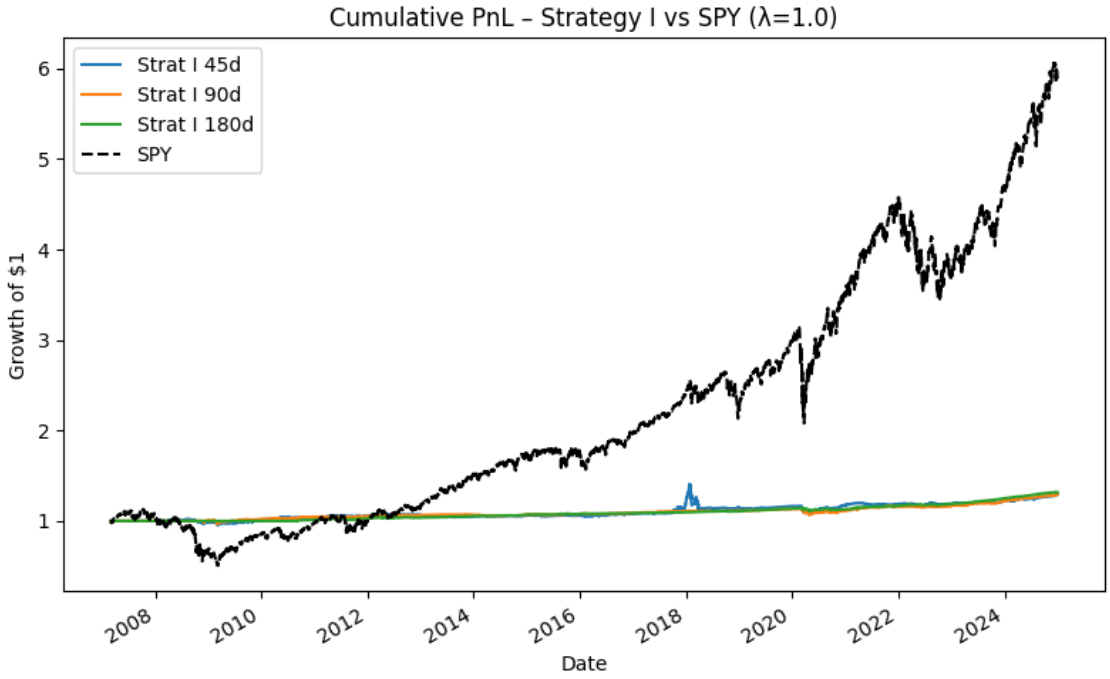


Figure 3: Cumulative Performance of Strategy I vs. SPY ($\lambda = 1.0$)

Figure 3 further demonstrates the impact of increasing risk aversion, with $\lambda = 1.0$. At this level, all three estimation windows produce remarkably similar and stable performance trajectories. While they no longer exhibit the potential for significant outperformance seen with lower λ values, they also avoid the dramatic drawdowns. However, this comes at the cost of underperforming the S&P 500 benchmark (black dashed line) over most of the analysis period.

These results suggest a clear trade-off between potential outperformance and risk management, with higher λ values providing more stable but lower returns. The optimal choice depends on investor preferences and market conditions, with lower λ values potentially advantageous during stable market environments and higher values during turbulent periods.

4.2 Performance Across Market Regimes

Next, we analyze how our strategies perform across different market regimes defined in our methodology section.

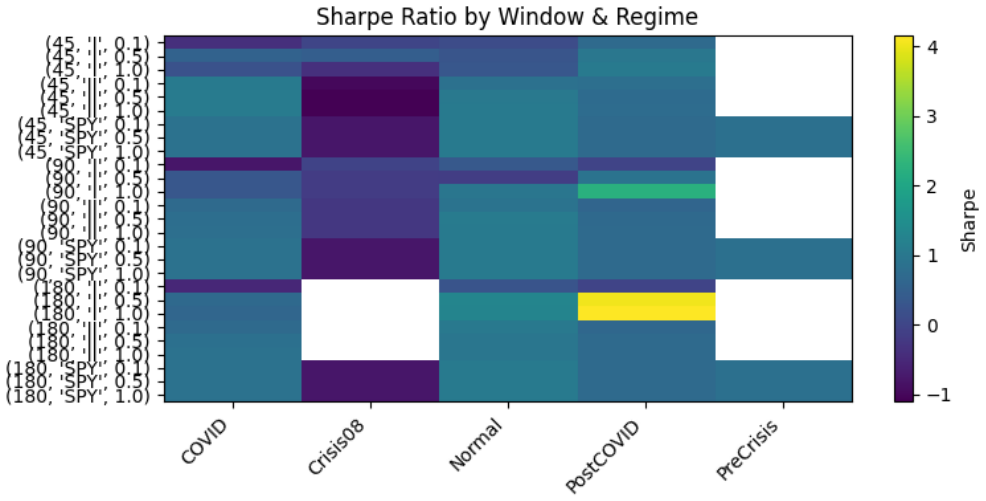


Figure 4: Sharpe Ratio by Window, Strategy, Lambda, and Market Regime

Figure 4 presents a heatmap of Sharpe ratios across different market regimes, strategies, estimation windows, and risk aversion parameters. The color intensity represents the Sharpe ratio value, with darker colors indicating lower values and brighter colors (particularly yellow) representing higher Sharpe ratios.

Several interesting patterns emerge from this visualization:

- **Crisis08 Period:** We observe generally lower Sharpe ratios (darker colors) across most strategy configurations during the 2008 financial crisis, which aligns with expectations given the extreme market volatility during this period. Notably, the 45-day window for Strategy I appears to perform relatively better than longer windows, suggesting that more responsive estimation methods have advantages during rapidly changing market conditions.

- **Normal Period:** During normal market conditions, we see moderately positive Sharpe ratios across most configurations, with 90-day windows demonstrating relatively consistent performance across different strategies and risk aversion parameters.

- **COVID Period:** The COVID market disruption shows interesting patterns, with some strategy configurations (particularly Strategy I with 180-day window and $\lambda = 1.0$) achieving remarkably high Sharpe ratios (bright yellow). This suggests that strategies with longer estimation windows and higher risk aversion performed exceptionally well during this specific crisis, possibly because they were less reactive to the initial market shock.

- **Post-COVID Period:** Performance in the post-COVID recovery shows mixed results, with several configurations showing white spaces, indicating potential numerical issues or extremely poor performance.

- **Pre-Crisis Period:** The pre-crisis period shows generally moderate performance across different configurations, with no clear standout performer.

The stark contrast in performance across different market regimes reinforces the importance of adapting estimation methods to market conditions. No single window length or risk parameter proves universally superior across all regimes, highlighting the potential value of a dynamic approach to parameter selection.

The heatmap also reveals that SPY (the benchmark) generally exhibits consistent Sharpe ratios across different regimes, which is visible in the consistent coloring of the SPY rows across different columns. This suggests that while our strategies can achieve superior risk-adjusted returns in certain regimes, they may also be more sensitive to regime changes compared to the benchmark.

4.3 Impact of Risk Aversion Parameter

The risk aversion parameter (λ) plays a crucial role in determining the risk-return profile of our strategies.

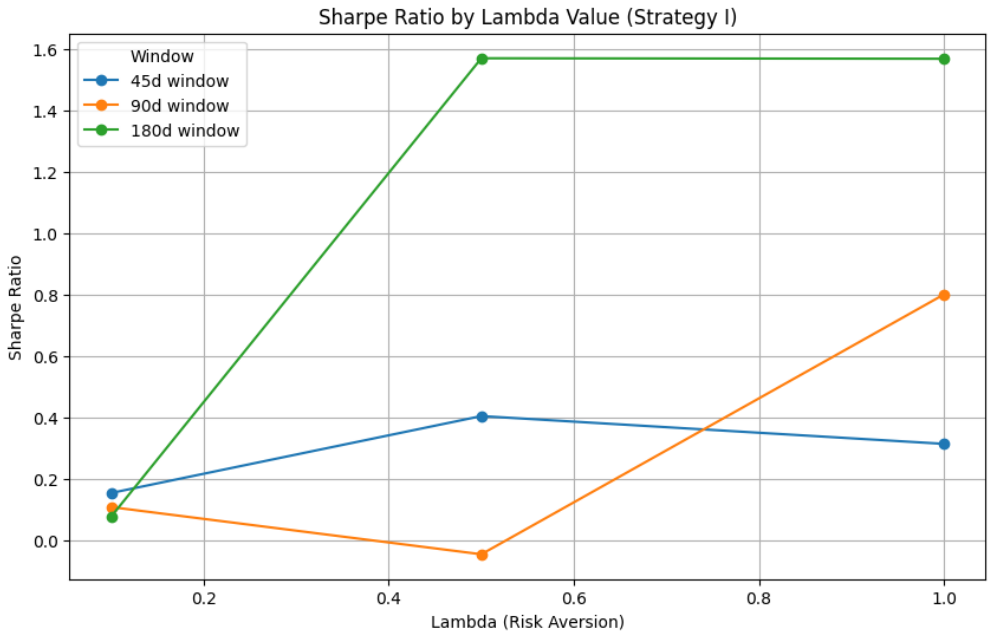


Figure 5: Sharpe Ratio by Lambda Value for Strategy I with Different Window Lengths

Figure 5 illustrates how the Sharpe ratio for Strategy I varies with the risk aversion parameter for different estimation windows. The results reveal important insights about the interaction between estimation windows and risk aversion:

- **180-day Window (Green Line):** The long-term estimation window shows the most dramatic sensitivity to the risk aversion parameter. At low λ values (around 0.1), it performs poorly with a Sharpe ratio near zero. However, as λ increases, its performance improves dramatically, achieving the highest Sharpe ratios among all windows at $\lambda = 0.5$ and maintaining this advantage through $\lambda = 1.0$. This suggests that longer estimation windows benefit substantially from increased risk control, likely because

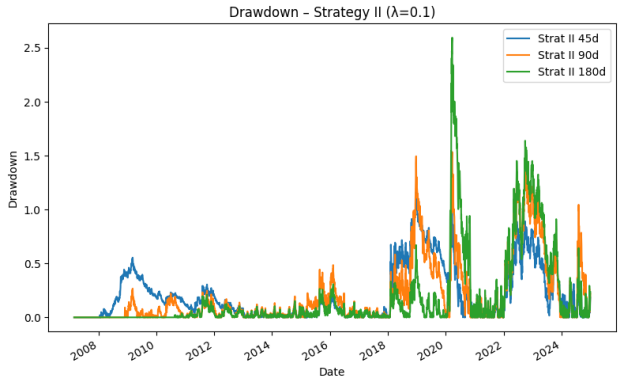
they provide more stable but potentially outdated estimates that require stronger risk constraints to perform optimally.

- **45-day Window (Blue Line):** The short-term estimation window shows moderate but relatively consistent performance across different λ values. It peaks at $\lambda = 0.5$ and then slightly declines at higher values. This more stable pattern suggests that short-term windows are less sensitive to risk aversion parameters, possibly because they already adapt quickly to changing market conditions.
- **90-day Window (Orange Line):** The mid-term estimation window displays the most complex relationship with λ. It shows poor performance at low and mid-range λ values, even dipping into negative territory at $\lambda = 0.5$, before recovering strongly at $\lambda = 1.0$. This U-shaped pattern suggests that the 90-day window requires either very low risk aversion (to capitalize on its market timing ability) or very high risk aversion (to control potential estimation errors).

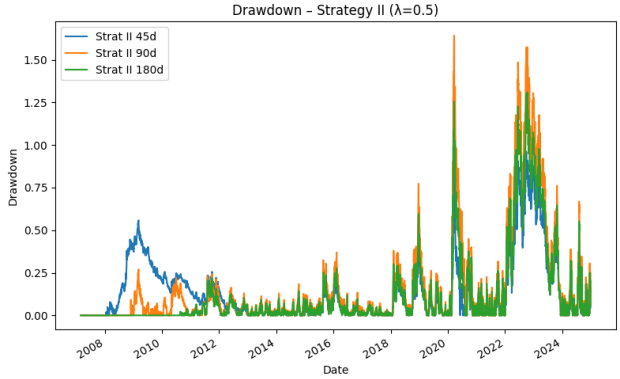
These results highlight the importance of properly matching the estimation window with an appropriate risk aversion parameter. The 180-day window clearly benefits from higher λ values (0.5-1.0), while the 45-day window performs reasonably well across a range of risk parameters. The 90-day window appears more challenging to optimize, requiring careful parameter selection based on market conditions.

4.4 Drawdown Analysis

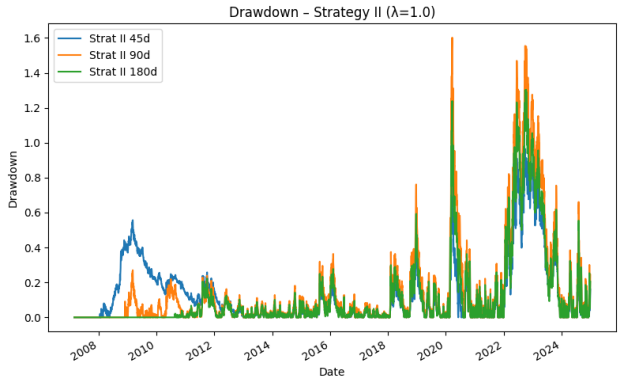
Examining drawdowns provides insights into the worst-case scenarios for our strategies and their ability to manage downside risk.



(a) Drawdown for Strategy II with $\lambda = 0.1$



(b) Drawdown for Strategy II with $\lambda = 0.5$



(c) Drawdown for Strategy II with $\lambda = 1.0$

Figure 6: Drawdown Profiles for Strategy II with Different Risk Aversion Parameters

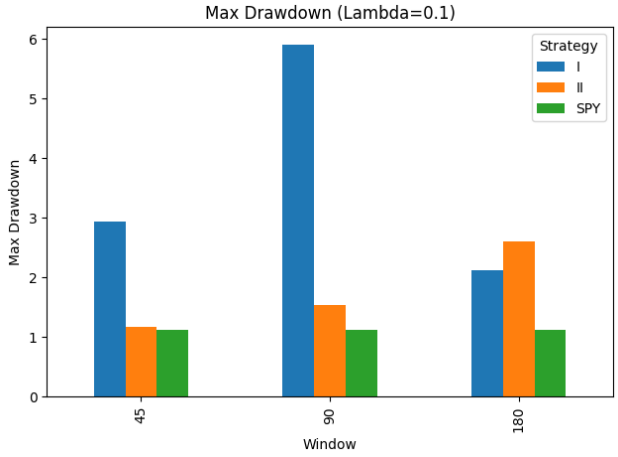
Figure 6 presents the drawdown profiles for Strategy II across different estimation windows and risk aversion parameters. Several key observations emerge:

- **Impact of λ on Maximum Drawdown:** As the risk aversion parameter increases from 0.1 to 1.0, we observe a general reduction in maximum drawdown across all window lengths. For instance, the 180-day window (green line) experiences a peak drawdown exceeding 2.5 at $\lambda = 0.1$ during the COVID period (2020), but this is reduced to approximately 1.6 at $\lambda = 0.5$ and further to around 1.2 at $\lambda = 1.0$. This

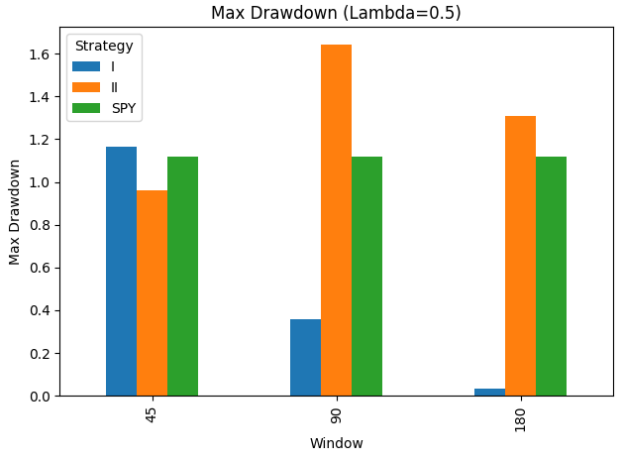
confirms that higher risk aversion effectively mitigates downside risk.

- **Temporal Patterns:** Major drawdown periods align with significant market events. We observe substantial drawdowns during the 2008-2010 financial crisis, a moderate drawdown period around 2016-2018, and the most severe drawdowns during the COVID-19 pandemic in 2020. The post-COVID period (2022-2024) also shows significant drawdowns, particularly for the 90-day and 180-day windows.

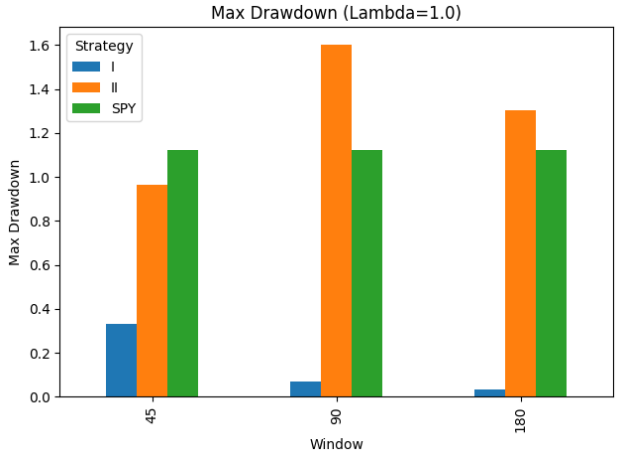
- **Window Length Comparison:** The 45-day window (blue line) typically shows earlier and sometimes deeper drawdowns during the initial phase of market crises (visible in 2008-2009), but often recovers more quickly. This pattern is consistent with the short-term window's greater responsiveness to market changes. In contrast, the 180-day window (green line) tends to experience more prolonged drawdown periods, particularly evident in the post-COVID period.



(a) Maximum Drawdown Comparison with $\lambda = 0.1$



(b) Maximum Drawdown Comparison with $\lambda = 0.5$



(c) Maximum Drawdown Comparison with $\lambda = 1.0$

Figure 7: Maximum Drawdown Comparison Across Strategies and Window Lengths

Figure 7 provides a comparative view of maximum drawdowns across both strategies and the SPY benchmark for different window lengths and risk aversion parameters. These bar charts reveal several important insights:

- **Strategy I vs. Strategy II:** At $\lambda = 0.1$, Strategy I (blue bars) experiences substantially higher maximum drawdowns than Strategy II (orange bars) across all window lengths, with particularly severe drawdowns for the 90-day window. This pattern dramatically reverses at higher λ values (0.5 and 1.0), where Strategy I achieves significantly lower maximum drawdowns than Strategy II, especially for the 90-day and 180-day windows.

- **Comparison to Benchmark:** The SPY benchmark (green bars) shows relatively consistent maximum drawdowns around 1.1 regardless of the comparison scenario. At $\lambda = 0.1$, both strategies generally experience higher maximum drawdowns than the benchmark. However, at $\lambda = 1.0$, Strategy I achieves substantially lower maximum drawdowns than SPY across all window lengths.

- **Optimal Window Length:** For Strategy I at higher λ values (0.5 and 1.0), the 180-day window consistently produces the lowest maximum drawdowns, followed by the 90-day window. This suggests that longer estimation windows, when combined with sufficient risk control, provide the best downside protection.

These drawdown analyses highlight the critical interplay between strategy design, estimation window, and risk aversion parameter. Strategy I with higher λ values and longer estimation windows provides the best downside protection, while Strategy II consistently experiences higher drawdowns regardless of parameter choices. This aligns with the fundamental design difference between the strategies, with Strategy I employing tighter beta constraints (-0.5 to 0.5) that limit market exposure.

5 Discussion and Recommendations

5.1 Optimal Estimator Selection

Based on our empirical results, we can derive recommendations for estimator selection under different market conditions:

- **Crisis Periods (2008 Financial Crisis, COVID-19):** For Strategy I during crisis periods, our results suggest that shorter-term estimators (45 days) combined with higher risk aversion parameters ($\lambda = 0.5$ or 1.0) provide the most effective combination. The 45-day window allows for quicker adaptation to rapidly changing market conditions, while the higher risk aversion mitigates excessive risk-taking. This is evidenced by the stronger performance of the 45-day window during the early phase of the 2008 crisis (as seen in Figure 1) and the lower drawdowns achieved with higher lambda values (Figure 7).

  For Strategy II, our drawdown analysis indicates that mid-term estimators (90 days) with moderate to high risk aversion ($\lambda = 0.5$ or 1.0) offer a better balance. The wider

beta constraints of Strategy II benefit from the additional stability provided by slightly longer estimation windows during turbulent periods.

- **Stable Market Conditions:** During stable market periods, our Sharpe ratio heatmap (Figure 4) shows that longer-term estimators (90-180 days) generally perform better for both strategies. For Strategy I, the 180-day window with high risk aversion ($\lambda = 1.0$) shows particularly strong risk-adjusted performance during normal market conditions. This combination benefits from the statistical robustness of longer estimation periods when market relationships are more stable.

  For Strategy II, the 90-day window with lower risk aversion ($\lambda = 0.1$) appears to offer the best balance during stable markets, allowing the strategy to capitalize on market opportunities while maintaining reasonable risk control.

- **Transitional Periods:** During transitions between market regimes, our analysis suggests that mid-term estimators (90 days) with moderate risk aversion ($\lambda = 0.5$) provide the best compromise. This combination balances adaptability to changing conditions with sufficient stability to avoid excessive turnover and whipsaw effects during volatile transitions.

5.2 Strategy Selection and Risk Management

Our analysis also yields insights into the appropriate choice of strategy and risk aversion parameter for different investor profiles:

- **Conservative Investors:** Strategy I with a higher risk aversion parameter ($\lambda = 0.5$ or 1.0) and a mid-to-long-term estimation window (90-180 days) offers the best risk-adjusted performance for investors prioritizing capital preservation. As shown in Figure 7, this combination achieves significantly lower maximum drawdowns compared to both Strategy II and the SPY benchmark, while still providing reasonable returns over the long term.

- **Aggressive Investors:** Strategy II with a lower risk aversion parameter ($\lambda = 0.1$) and a dynamic approach to window selection—using longer estimation windows (180 days) during stable markets and shorter windows (45-90 days) during periods of market stress—provides higher potential returns at the cost of increased volatility. This approach leverages the strategy's wider beta constraints to capitalize on market opportunities while adapting to changing conditions.

- **Balanced Approach:** For investors seeking a middle ground, Strategy I with moderate risk aversion ($\lambda = 0.5$) and a 45-day window offers an attractive compromise. As shown in Figure 2, this combination achieves competitive long-term performance while maintaining reasonable downside protection, making it suitable for investors with moderate risk tolerance.

5.3 Practical Implementation Considerations

Several practical considerations emerge from our analysis that would be valuable for portfolio managers implementing these strategies:

- **Regime Detection:** Implementing a regime-switching approach requires reliable methods for detecting changes in market conditions. Simple approaches based on market volatility (such as the VIX index) or moving average crossovers could provide useful signals for adjusting estimation windows and risk parameters in real-time.
- **Computational Efficiency:** Strategy I (using convex optimization) is generally more computationally efficient and robust than Strategy II (requiring nonlinear optimization), making it potentially more suitable for practical implementation, especially when frequent rebalancing is required. Our implementation using CVXPY for Strategy I demonstrated greater numerical stability compared to the SLSQP method used for Strategy II.
- **Transaction Costs:** While not explicitly modeled in our analysis, strategies using shorter estimation windows typically generate higher turnover, potentially increasing transaction costs. Our examination of weekly weights (not shown in this report) indicates that the 45-day window generally results in more frequent and larger position changes compared to the 90-day and 180-day windows. Practical implementations should incorporate transaction cost models and possibly include turnover constraints in the optimization.
- **Hybrid Approaches:** Our results suggest potential benefits from hybrid approaches that use different estimation windows for different aspects of the model. For example, using longer windows (180 days) for covariance estimation to enhance stability, while employing shorter windows (45 days) for expected return estimation to capture more recent trends, could provide an optimal balance of responsiveness and robustness.

6 Conclusion

In this project, we have analyzed two factor-based long-short allocation strategies with different beta constraints, focusing on the impact of estimation windows and risk aversion parameters on performance across various market regimes.

Our analysis reveals that estimation window selection critically impacts strategy performance. Short-term estimators (45 days) excel during market stress when combined with higher risk aversion, allowing rapid adaptation while maintaining risk control. Conversely, longer-term windows (90-180 days) outperform during stable conditions, benefiting from reduced estimation noise and statistical robustness.

The risk aversion parameter (λ) significantly influences strategy behavior. For Strategy I with tight beta constraints, higher λ values substantially improve performance, particularly with longer estimation windows. Strategy II with wider beta constraints performs best with moderate λ values that balance return potential with risk management.

Comparing the strategies, we find Strategy I achieves superior downside protection with higher λ values, especially during crisis periods, while Strategy II allows greater market exposure that can enhance returns during favorable conditions but increases vulnerability during downturns. This demonstrates how beta constraints fundamentally shape portfolio behavior and risk characteristics.

Our findings emphasize that no single parameter combination proves universally optimal. A dynamic approach that adapts estimation techniques to changing market conditions is likely to yield the best long-term results. Practical implementations should consider regime detection methodologies, computational efficiency tradeoffs, and transaction cost implications.

Future research directions include incorporating transaction costs, developing sophisticated regime-detection methods, and testing hybrid approaches that use different estimation windows for various model components. These enhancements could further refine the practical application of factor-based allocation strategies across diverse market environments.