

# A Dual-Layer, Consensus-Based, Multi-Factor Framework for Market-Neutral Long/Short Equity Investing

\*A project by Queen’s University Algorithmic Network and Trading Team (QUANTT)

1<sup>st</sup> Benjamin Gorenc  
*Quantitative Analyst*  
QUANTT  
Kingston, ON  
B.Gorenc@queensu.ca

2<sup>nd</sup> Gabriel Soler  
*Quantitative Analyst*  
QUANTT  
Kingston, ON  
Gabe.Soler@queensu.ca

3<sup>rd</sup> Andrew Farkouh  
*Quantitative Analyst*  
QUANTT  
Kingston, ON  
22wmk22@queensu.ca

**Abstract**—This paper presents a systematic, market-neutral long/short equity strategy operating across a universe of 399 U.S.-listed equities and ADRs. Alpha is generated via a composite scoring engine that normalizes and dynamically weights five fundamental and price-based factors: Momentum, Value, Quality, Low Volatility, and Size. To solve for optimal position sizing, we introduce the Consensus-Based Optimizer (CBO)—a novel distributed gradient-descent algorithm. The CBO models equities as independent agents on a fully connected network, converging on a consensus allocation that actively penalizes portfolio variance, L2 concentration, and turnover. Sector-level capital is allocated using a Black-Litterman enhanced mean-variance framework with Ledoit-Wolf covariance shrinkage, while strict sleeve-level beta neutralization ( $\pm 0.02$ ) is enforced via iterative short-leg rescaling. Backtested from March 2025 to March 2026, the strategy delivers a 19.73% annualized return, a Sharpe ratio of 2.51, and a maximum drawdown of just  $-2.22\%$ . These results validate that the CBO’s multi-agent network approach effectively isolates robust, market-independent alpha with materially lower downside risk than the SPY benchmark.

## I. INTRODUCTION

### A. Background

Long/short equity strategies have been the backbone of systematic hedge fund investing since the work of Fama and French [1], who demonstrated that cross-sectional return dispersion across stocks is persistent, predictable, and attributable to identifiable risk factors. A portfolio that goes long on stocks with favorable factor exposures while simultaneously shorting those with unfavorable ones can, in principle, harvest this dispersion as market-neutral alpha — returns that are structurally uncorrelated with broad equity market movements. However, the practical realization of such strategies requires solving three non-trivial problems in sequence: identifying reliable cross-sectional signals, allocating capital across sectors in a risk-aware manner, and sizing individual positions within each sector without introducing unintended directional exposure.

Traditional implementations address position sizing through mean-variance optimization (MVO) [2] or equal-weighting. MVO is theoretically elegant but notoriously sensitive to

estimation error in expected returns, producing unstable, concentrated allocations that perform poorly out of sample [3]. Equal-weighting ignores the information content of factor scores entirely, treating a stock with a composite score of 3.0 identically to one scoring 0.9. Neither approach is satisfactory for a strategy that depends precisely on the relative strength of factor signals within a sector.

This paper introduces a systematic pipeline developed by the Queen’s University Algorithmic Network and Trading Team (QUANTT), grounded in academic research on fully connected graphs, network topology, and decentralized consensus mechanisms. The proposed model addresses each of the core challenges in mean-variance optimization (MVO), with particular emphasis on a novel position-sizing algorithm. The Consensus-Based Optimizer (CBO) replaces both MVO and equal-weighting with a distributed, multi-agent consensus mechanism, offering a more robust and interpretable framework for cross-sectional portfolio construction.

### B. Problem Definition

The central problem is the construction of a beta-neutral, multi-factor, long/short equity portfolio over a universe of US-listed equities. Formally, the goal is to find a weight vector  $\mathbf{w} \in \mathbb{R}^n$  for each sector containing  $n$  stocks, such that:

$$\min_{\mathbf{w}} \mathbf{w}^\top \Sigma \mathbf{w} + \lambda_\beta (\boldsymbol{\beta}^\top \mathbf{w})^2 + \lambda_2 \|\mathbf{w}\|^2 + \lambda_\tau \|\mathbf{w} - \mathbf{w}_{\text{prev}}\|_1 - \boldsymbol{\alpha}^\top \mathbf{w} \quad (1)$$

where  $\mathbf{w}$  represents the vector of target portfolio weights,  $\Sigma$  is the estimated covariance matrix of asset returns, and  $\boldsymbol{\alpha}$  denotes the vector of expected returns (or cross-sectional factor signals). The objective function consists of five distinct components designed to balance max returns with risk and implementation constraints:

- $\mathbf{w}^\top \Sigma \mathbf{w}$ : The standard portfolio variance, serving as the risk penalty.

- $\lambda_\beta(\beta^\top \mathbf{w})^2$ : A penalty on the portfolio’s net market exposure, mathematically enforcing the strategy’s beta-neutral mandate.
- $\lambda_2\|\mathbf{w}\|^2$ : An  $L_2$  (Ridge) regularization term that severely penalizes concentrated positions. This directly mitigates the instability typical of standard mean-variance optimization, promoting a diversified and robust allocation.
- $\lambda_\tau\|\mathbf{w} - \mathbf{w}_{\text{prev}}\|_1$ : An  $L_1$  penalty on the absolute change in weights from the previous period ( $\mathbf{w}_{\text{prev}}$ ), which acts as a proxy for turnover, slippage, and transaction costs.
- $-\alpha^\top \mathbf{w}$ : The expected portfolio return, subtracted here to align return maximization with the overarching minimization framework.

The non-negative coefficients  $\lambda_\beta$ ,  $\lambda_2$ , and  $\lambda_\tau$  act as hyper-parameters that dictate the relative importance of beta neutrality, diversification, and turnover reduction against the core mean-variance trade-off.

### C. Contributions

This paper makes three primary contributions:

- 1) A **five-factor scoring engine** with sector-specific weights, combining Momentum, Value, Quality, Low Volatility, and Size signals constructed from adjusted price and fundamental data, normalized cross-sectionally via winsorization and z-scoring.
- 2) **The Consensus-Based Optimizer (CBO)**, a novel distributed gradient-descent algorithm in which each stock in a sleeve acts as an independent agent on a fully connected network, converging to a consensus weight vector through 400 iterations of local gradient updates blended with population-level averaging.
- 3) **An end-to-end portfolio construction pipeline** integrating Black-Litterman enhanced sector allocation with Ledoit-Wolf covariance shrinkage, risk-parity floors, and sleeve-level beta neutralization, backtested over March 2025 to March 2026 against a SPY total-return benchmark.

Backtesting demonstrates a 19.73% annualized return, a Sharpe ratio of 2.51, and a maximum drawdown of only  $-2.22\%$ , validating the effectiveness of the CBO’s risk-aware consensus mechanism for generating consistent, market-independent alpha.

## II. DATA SOURCES & UNIVERSE CONSTRUCTION

### A. Universe Screening

The investment universe is constructed by screening U.S.-listed equities across 11 Global Industry Classification Standard (GICS) sectors using the Yahoo Finance Equity Query API. For each sector, the top 50 stocks by market capitalization are retrieved. After which, sector specific minimum market capitalization thresholds are applied to remove small cap companies. Table I summarizes the market capitalization thresholds applied to each sector.

After screening, duplicate listings are removed because the same company may appear under multiple ticker symbols; the

TABLE I  
SECTOR-SPECIFIC UNIVERSE SCREENING PARAMETERS.

Sector	Min. Cap (USD)	Rationale
Technology	\$100M	Exclude speculative micro-caps; data integrity
Healthcare	\$50M	Include mid-cap biotech with sufficient liquidity
Financial Services	\$75M	Exclude micro-cap financials; valuation data quality
Energy	\$30M	Include smaller E&P names for value factor breadth
All Other Sectors	\$20M	Default minimum; captures adequate universe depth

highest market capitalization entry for each unique company is retained. The resulting raw universe spans approximately 400 unique U.S. and international American Depositary Receipt (ADR) equities. Price data is then downloaded via yfinance and equities with less than 75% of data reported are removed. The active universe used for factor computation comprises 399 tickers across all 11 sectors. Table II reports the breakdown per-sector of the final active universe.

TABLE II  
FINAL ACTIVE UNIVERSE COMPOSITION BY SECTOR.

Sector	Some Representative Names	$n$
Technology	NVDA, AAPL, MSFT, TSM, AVGO	43
Healthcare	LLY, JNJ, ABBV, AZN, MRK	39
Financial Svcs	BRK-B, JPM, V, BAC, MA	30
Consumer Cyc.	AMZN, TSLA, HD, MCD, BKNG	37
Comm. Svcs	GOOGL, META, NFLX, TMUS, VZ	34
Industrials	GE, CAT, RTX, GEV, BA	33
Consumer Def.	WMT, COST, PG, KO, PM	37
Energy	XOM, CVX, TTE, COP, ENB	41
Utilities	NEE, CEG, SO, DUK, VST	40
Real Estate	WELL, PLD, EQIX, AMT, SPG	33
Materials	LIN, BHPLF, SCCO, NEM, FCX	32
<b>Total</b>		<b>399</b>

The universe includes a meaningful proportion of international names listed as ADRs, such as TSM (Taiwan Semiconductor), RHHBF (Roche), NVSEF (Novartis), TCTZF (Tencent), and TOYOF (Toyota). Including global names enhances factor diversification, particularly for the value and low-volatility factors, but introduces currency rates.

### B. Price Data

Adjusted daily closing prices are downloaded from Yahoo Finance for all 399 active tickers over the period January 1, 2020, through January 1, 2026, yielding a full price matrix of 1,508 trading days. Data is cached locally in Parquet format to avoid repeated API calls. For factor computation, a rolling

lookback window of 756 trading days (approximately three years) is applied, ending at the most recent available date, December 31st, 2025. This window is chosen to balance the need for a sufficiently long estimation period for momentum and low-volatility factors against the desire to keep factor signals responsive to current market conditions.

Market returns for beta computation are proxied by the SPY ETF (SPDR S&P 500 ETF Trust), downloaded over the same date range and reindexed to the universe price calendar. Individual stock betas are estimated by regressing daily stock returns on SPY returns over the 756-day lookback window, given by:

$$\beta_i = \frac{\text{Cov}(r_i, r_m)}{\text{Var}(r_m)} \quad (2)$$

For tickers with fewer than 60 observations, beta is set to 1.0 as a conservative default. Across the 399-ticker universe, the estimated betas range from  $-0.456$  to  $2.348$  with a cross-sectional mean of  $0.602$ . The below-unity mean reflects the significant weight of international ADRs, utilities, and consumer defensive names in the universe, which exhibit lower co-movement with the U.S. equity market.

### C. Fundamental Data

Fundamental data required for the value and quality factors is fetched via the Yahoo Finance Ticker API using a multi-threaded parallel download across 10 concurrent workers. The following nine fields are retrieved per ticker: book value per share, shares outstanding, market capitalization, forward EPS, net income, total revenue, total debt, total cash, and EBITDA. Derived fields book equity, total earnings, and total equity are computed from these base fields.

All fundamental values are held constant as a time series across the full price date index. An important limitation of this approach is that the fundamental data reflects point-in-time live values rather than historical quarterly snapshots. This means that the fundamental-based factors (value, quality) are computed using current financials applied retroactively across the backtest window, introducing a form of look-ahead bias for the historical backtest. As a result, the backtest should be interpreted as a live signal demonstration rather than a fully rigorous historical simulation.

## III. FACTOR CONSTRUCTION & COMPOSITE SCORING

The strategy implements five systematic factors drawn from the empirical asset pricing literature: Momentum (M), Value (V), Quality (Q), Low Volatility (LV), and Size (S). Each factor is constructed from price and fundamental data, normalized cross-sectionally via winsorization and z-scoring, and combined into a single composite score per stock.

### A. Momentum Factor (M)

The momentum factor captures the tendency of recent outperforming stocks to continue outperforming in the near term [6], [7]. Two momentum signals are computed, normalized, and averaged. The first is a standard 12-month minus

one-month (skip-month) return, which avoids the short-term reversal effect:

$$M_1(t) = \frac{P(t-21)}{P(t-252)} - 1 \quad (3)$$

where  $P(t)$  denotes the adjusted closing price at time  $t$ , and lags are expressed in trading days (21 days  $\approx$  1 month, 252 days  $\approx$  one year). The one-month skip ensures the signal reflects intermediate-horizon momentum rather than short-term noise. The second signal is a six-month price return:

$$M_2(t) = \frac{P(t)}{P(t-126)} - 1 \quad (4)$$

where 126 trading days  $\approx$  six months. The composite momentum factor is the simple average of both signals. Note that both momentum scores are individually normalized before an average is computed:

$$M(t) = \frac{1}{2} [M_1(t) + M_2(t)] \quad (5)$$

Averaging the two horizons smooths idiosyncratic noise in either measure and has been shown to produce more stable momentum signals than either measure alone [8]. In the current portfolio as of December 2025, momentum is the primary driver of signals in Technology (weight 35%) and Consumer Cyclical (weight 30%), consistent with the higher return dispersion in those sectors.

### B. Value Factor (V)

The value factor identifies stocks trading at low prices relative to their fundamental worth [9]. Four valuation sub-signals are computed and averaged to form a robust composite value score:

$$V_1 = \frac{\text{Book Equity}}{\text{Market Cap}} \quad (6)$$

$$V_2 = \frac{\text{Earnings}}{\text{Price}} \quad (7)$$

$$V_3 = \frac{\text{EBITDA}}{\text{Market Cap} + \text{Debt} - \text{Cash}} \quad (8)$$

$$V_4 = \frac{\text{Book Value}}{\text{Price}} \quad (9)$$

All four sub-signals are defined such that higher values indicate cheaper, more attractively valued stocks. Infinite and undefined values are replaced with NaN and excluded from the average. All four value formulas are independently normalized prior to average computation.

### C. Quality Factor (Q)

The quality factor selects stocks with superior profitability, earnings stability, and risk-adjusted returns [11]. Four sub-

signals are computed:

$$Q_1 = \frac{\text{Net Income}}{\text{Book Equity}} \quad (10)$$

$$Q_2 = \frac{\text{Net Income}}{\text{Revenue}} \quad (11)$$

$$Q_3(t) = \frac{\mu_{60}(t)}{\sigma_{60}(t)} \quad (12)$$

$$Q_4 = \frac{\mu_{E,8}}{\sigma_{E,8}} \quad (13)$$

Equation 10 is return on equity (ROE). Equation 11 is profit margin. Equation 12 is the 60-day rolling return-to-risk ratio of daily returns. Equation 13 captures earnings stability, defined as the inverse of the coefficient of variation of earnings over a rolling 8-period window. All four quality formulas are independently normalized prior to average computation.

#### D. Low Volatility Factor (LV)

The low volatility factor exploits the empirical anomaly where lower-risk stocks tend to yield higher risk-adjusted returns than high-beta, highly volatile stocks [12]. This factor combines two sub-signals based on a 60-day rolling window:

$$LV_1(t) = \frac{1}{\sigma_{60}(t)} \quad (14)$$

$$LV_2(t) = \frac{1}{\sigma_{\epsilon,60}(t)} \quad (15)$$

Equation 14 represents total inverse volatility. Equation 15 isolates idiosyncratic inverse volatility, where  $\sigma_{\epsilon,60}(t)$  is the rolling standard deviation of the residuals ( $\epsilon$ ) derived from regressing the asset's daily returns against the market return.

#### E. Size Factor (S)

The size factor targets the small-cap premium [10] by identifying firms with smaller market capitalizations. To reduce the extreme positive skewness inherent in market capitalization distributions and allow for stable cross-sectional z-scoring, the signal is defined as the negative natural logarithm of the firm's market capitalization:

$$S(t) = -\ln(\text{Market Cap}) \quad (16)$$

This mathematically enforces a small-minus-big (SMB) tilt, assigning higher factor scores to smaller firms.

### IV. CONSENSUS-BASED OPTIMIZER (CBO)

To determine optimal position sizing within each sector sleeve, this strategy employs a novel Consensus-Based Optimizer (CBO). The CBO models each stock as an independent agent operating on a fully connected network, using a distributed gradient-descent mechanism to iteratively update weights based on local information and neighbor communication. The algorithm converges when  $\|\mathbf{w}^{(t)} - \mathbf{w}^{(t-1)}\|_2 < \epsilon$ , at which point all agents have reached consensus and no further weight updates occur.

#### A. Alpha and Covariance Estimation

The expected return vector,  $\alpha$ , is constructed using the composite factor scores of the selected equities. To mitigate the impact of outliers, raw factor scores ( $s_i$ ) are normalized cross-sectionally using a z-score clipped between  $-3.0$  and  $3.0$ :

$$z_i = \text{clip}\left(\frac{s_i - \mu_s}{\sigma_s}, -3.0, 3.0\right) \quad (17)$$

The alpha vector is then generated by passing the normalized scores through a rectified linear activation, scaled by a base signal magnitude ( $\mu_{\text{base}} = 1.0$ ) and floored by a small epsilon factor ( $\epsilon = 1 \times 10^{-6}$ ) to ensure all alphas remain strictly positive:

$$\alpha_i = \mu_{\text{base}} (\max(z_i, 0) + \epsilon) \quad (18)$$

The covariance matrix,  $\Sigma$ , is estimated using a 126-day rolling window of daily returns. To ensure numerical stability and invertibility, the sample covariance is regularized via Ledoit-Wolf shrinkage ( $\Sigma_{\text{LW}}$ ) [5], which is further blended with its own diagonal using an additional shrinkage intensity parameter ( $\delta = 0.20$ ):

$$\Sigma = (1 - \delta)\Sigma_{\text{LW}} + \delta \text{diag}(\Sigma_{\text{LW}}) + \epsilon I \quad (19)$$

This two-stage regularization suppresses off-diagonal comovements while preserving asset-level variance estimates, causing the optimizer to more heavily penalize idiosyncratic risk relative to a standard Ledoit-Wolf matrix alone.

#### B. Distributed Gradient Descent

The optimization process iterates 400 times ( $N_{\text{steps}} = 400$ ) with a learning rate step size ( $\eta$ ) of 0.04. The distributed gradient-descent framework follows the multi-agent subgradient approach of Nedić and Ozdaglar [13]. At each iteration  $t$ , the algorithm calculates the local gradient of the objective function for each agent  $i$ , which incorporates the four primary penalty terms:

$$\begin{aligned} \nabla f(\mathbf{w}_i) = & -\alpha + 2\lambda_{\text{risk}}\Sigma\mathbf{w}_i + 2\lambda_2\mathbf{w}_i \\ & + 2\lambda_\beta(\beta^\top \mathbf{w}_i)\beta + \lambda_\tau \text{sign}(\mathbf{w}_i - \mathbf{w}_{\text{prev}}) \end{aligned} \quad (20)$$

The hyperparameters are: Risk Aversion ( $\lambda_{\text{risk}} = 8.0$ ), Beta Penalty ( $\lambda_\beta = 15.0$ ), Concentration Penalty ( $\lambda_2 = 0.05$ ), and Turnover Penalty ( $\lambda_\tau = 0.02$ ).

Following the local gradient step, the weights are updated by blending the agent's local gradient descent step ( $\gamma_1 = 0.75$ ), the network consensus mixing matrix ( $\gamma_2 = 0.20$ ), and the global population mean ( $\gamma_3 = 0.05$ ):

$$\mathbf{w}_i^{(t+1)} = \gamma_1 \left( \mathbf{w}_i^{(t)} - \eta \nabla f(\mathbf{w}_i^{(t)}) \right) + \gamma_2 \sum_j M_{ij} \mathbf{w}_j^{(t)} + \gamma_3 \bar{\mathbf{w}}^{(t)} \quad (21)$$

where the consensus mixing matrix  $M = I - \epsilon_{\text{mix}} L$  with  $\epsilon_{\text{mix}} = [\text{value}]$  and  $L$  representing the Graph Laplacian matrix [14] allows agents to share information across a fully connected adjacency network topology. To ensure the update is a convex combination of the three signals,  $\gamma_1 + \gamma_2 + \gamma_3 = 1$ .

### C. Beta Neutralization

After consensus is reached, a two-pass sleeve-level beta neutralization procedure is used to ensure  $|\beta_p| < 0.02$ . On the first pass, each sleeve's weight vector is independently neutralized. On the second pass, a portfolio-level polish is applied across all sleeves sequentially. In each pass, the procedure iterates up to 30 steps. At each step, the net portfolio beta is computed:

$$\beta_p = \sum_i w_i \hat{\beta}_i \quad (22)$$

if  $|\beta_p| > 0.02$ , the procedure blends the long and short legs towards a beta-offsetting target distribution with step size  $\eta = 0.20$ .

$$w_L^{(k+1)} = (1 - \eta) w_L^{(k)} + \eta \tilde{w}_L \quad (23)$$

$$w_S^{(k+1)} = (1 - \eta) w_S^{(k)} + \eta \tilde{w}_S \quad (24)$$

where  $\tilde{w}_L \propto 1/\hat{\beta}_L$  and  $\tilde{w}_S \propto \hat{\beta}_S$  when  $\beta_p > 0$ , and reversed when  $\beta_p < 0$ . Each leg is re-normalized and re-capped at 20% after each step. The procedure terminates when  $|\beta_p| \leq 0.02$ .

## V. SECTOR ALLOCATION VIA MEAN-VARIANCE OPTIMIZATION

Sector-level capital allocation is determined by a Mean-Variance Optimization (MVO) model developed by Markowitz [2], enhanced with three modifications that address the well-known failure modes of the traditional MVO: Black-Litterman implied equilibrium returns [4] replace noisy historical means, Ledoit-Wolf shrinkage [5] stabilizes the covariance matrix, and risk-parity floors ensure minimum exposure to low-volatility defensive sectors. The optimizer determines the vector of sector sleeve weights  $w = (w_1, \dots, w_N)$  that maximizes the portfolio Sharpe ratio subject to sector bounds.

*Note on sector coverage:* The investment universe spans 11 GICS sectors (Table II), but the MVO sleeve allocation covers 8 sectors: Communication Services, Consumer Defensive, Energy, Financial Services, Healthcare, Industrials, Technology, and Utilities. Consumer Cyclical, Real Estate, and Materials are included in the universe for factor scoring but are excluded from portfolio construction in the current implementation.

### A. Covariance Estimation via Ledoit-Wolf Shrinkage

The covariance matrix is the central input to the optimizer, from which both portfolio variance and Black-Litterman implied returns are derived. For  $N$  sectors, the covariance matrix  $\Sigma$  is a symmetric  $N \times N$  matrix where each diagonal entry  $\sigma_i$  is the annualized variance of sector  $i$  and each off-diagonal entry  $\sigma_{ij}$  captures the co-movement between sectors  $i$  and  $j$ . Portfolio variance is given by the quadratic form:

$$\text{Var}(\mathbf{w}) = \mathbf{w}^T \Sigma \mathbf{w} \quad (25)$$

which accounts for all pairwise covariance terms simultaneously. Each sector's return is calculated as a weighted average of the returns of the individual stocks that belong to that sector, where the weights are determined by the CBO-assigned position sizes.

With eight sector sleeves and 756 observations, the sample covariance matrix contains 28 unique off-diagonal entries, some of which will reflect statistical noise rather than true structural co-movement. Raw MVO is known to treat all entries as equally reliable, aggressively exploiting false correlations by producing unstable, concentrated allocations; a phenomenon described by Michaud [3] as "error maximization." To address this, Ledoit-Wolf shrinkage is applied, blending the noisy sample matrix toward a structured target:

$$\Sigma_{LW} = (1 - \alpha) \Sigma_{\text{sample}} + \alpha \Sigma_{\text{structured}} \quad (26)$$

where  $\alpha$  is the shrinkage coefficient solved analytically to minimize expected estimation error, and  $\Sigma_{\text{structured}}$  is a constant-correlation target matrix with average variance on the diagonal. The Ledoit-Wolf estimator is guaranteed to produce a positive semi-definite matrix and requires no user-specified tuning parameter, making it well-suited to the relatively small cross-section of eight sectors [5].

### B. Black-Litterman Implied Equilibrium Returns

Raw MVO using historical mean returns as the expected return input produces allocations that reflect backward-looking performance rather than forward-looking risk compensation. Over the 2022–2025 estimation window, the AI-driven Technology rally and relative Energy underperformance would cause the optimizer to concentrate near the Technology cap and floor Energy at the minimum, an outcome that chases performance rather than ensuring proper diversification. The Black-Litterman model [4] resolves this by replacing historical means with implied equilibrium returns: the expected returns that a rational investor holding an equal-weight portfolio would require, given the covariance structure. These are derived as:

$$\Pi = \lambda \Sigma w_{\text{eq}} \quad (27)$$

where  $\Pi$  is the vector of implied equilibrium excess returns,  $\lambda = 2.5$  is the risk aversion scalar,  $\Sigma$  is the Ledoit-Wolf covariance matrix, and  $w_{\text{eq}} = (1/N, \dots, 1/N)$  is the equal-weight prior vector. The matrix-vector product  $\Sigma w_{\text{eq}}$  computes the risk contribution of each sector to the equal-weight portfolio. This describes how much each sector contributes to total portfolio variance when all sectors are held equally. Multiplying by  $\lambda$  scales this risk contribution into an implied return premium: sectors with higher individual variance or higher correlation to other sectors receive higher implied returns. A critical consequence of this formulation is that high-volatility sectors such as Energy and Technology receive higher implied returns than their recent historical performance suggests, while low-volatility sectors such as Consumer Defensive and Utilities receive more modest implied returns. This produces a more diversified optimizer output that does not mechanically chase recent winners. The present implementation uses no active views (the pure equilibrium prior is applied) though the Black-Litterman framework supports the addition of directional sector views through the standard  $P$ ,  $Q$ , and  $\Omega$  matrices, which is an area for future work.

### C. Max-Sharpe Optimization with Risk-Parity Floors

The optimizer maximizes the portfolio Sharpe ratio subject to per-sector lower bounds (floors) and a uniform upper bound (sector cap). The formal optimization problem is:

$$\text{Maximize: } S(w) = (w^\top \Pi - r_f) / \sqrt{w^\top \Sigma w} \quad (28)$$

$$\text{Subject to: } \sum w_i = 1, \quad \text{floor}_i \leq w_i \leq 0.35 \quad (29)$$

where  $r_f = 0.04$  is the annualized risk-free rate. The objective is nonlinear due to the square root in the denominator and has no closed-form solution; it is solved numerically using Sequential Least Squares Programming (SLSQP) with a convergence tolerance of  $\text{ftol} = 10^{-12}$ . The optimizer is warm-started at the risk-parity floors with remaining budget allocated proportionally to BL implied returns, ensuring feasibility from the first iteration and directing the search toward the likely optimal region.

Even with Black-Litterman expected returns, the Sharpe optimizer exhibits a systematic bias against low-volatility sectors. A sector such as Utilities has a modest BL implied return and contributes to portfolio variance; the optimizer may therefore reduce its allocation to lower the variance denominator, failing to fully credit the diversification benefit arising from Utilities' low correlation with Technology and Financial Services. To counteract this, risk-parity floors are applied per sector, set inversely proportional to volatility:

$$\text{Floor}_i = h \times (\bar{\sigma} / \sigma_i) \quad (30)$$

where  $h = 0.05$  is the universal hard floor,  $\bar{\sigma}$  is the average annualized volatility across all sectors, and  $\sigma_i$  is the annualized volatility of sector  $i$ , computed as the square root of the corresponding diagonal entry of  $\Sigma$ . This formula assigns a floor above  $h$  to low-volatility sectors (ratio  $> 1$ ) and clips high-volatility sectors at  $h$  (ratio  $< 1$ ). An Energy-specific override raises the floor to 7.0% to ensure meaningful commodity exposure independent of the formula result. If the sum of all floors exceeds 100%, floors are scaled proportionally with a 5% headroom buffer to preserve feasibility.

### D. Optimization Results

Table 5 reports the risk-parity floors and resulting sleeve weights produced by the optimizer over the 756-day estimation window ending December 31, 2025. Consumer Defensive receives the highest formula-derived floor (7.34%) due to its lowest annualized volatility (12.99%) among the eight sectors, consistent with its bond-proxy character. Technology and Financial Services are floored at the hard floor minimum of 5.00% as high-volatility sectors whose formula result falls below  $h$ .

Financial Services is allocated the full 35% sector cap, indicating that the optimizer sought a higher allocation but was constrained. This reflects the high BL implied return assigned to Financial Services — a consequence of its substantial covariance with the broader market, which the Black-Litterman framework interprets as requiring a higher equilibrium return

TABLE III  
MVO SECTOR ALLOCATION RESULTS. ANNUALIZED VOLATILITIES, RISK-PARITY FLOORS, AND FINAL SLEEVE WEIGHTS. ROLLING 3-YEAR ESTIMATION WINDOW ENDING DECEMBER 31, 2025.

Sector	Ann. Vol.	Floor	Sleeve Wt.
Comm. Services	18.12%	5.26%	5.26%
Consumer Def.	12.99%	7.34%	7.34%
Energy	23.08%	7.00%	7.00%
Financial Svcs	21.84%	5.00%	35.00%
Healthcare	16.52%	5.77%	5.77%
Industrials	17.53%	5.44%	5.44%
Technology	25.36%	5.00%	28.61%
Utilities	17.05%	5.59%	5.59%

premium. Technology is the only sector allocated freely between its floor and cap at 28.61%, representing the model's genuine unconstrained recommendation. All remaining sectors sit at their risk-parity floors, meaning the optimizer would have reduced their allocations further in the absence of constraints; they are held for diversification and sector exposure rather than pure Sharpe contribution.

The resulting allocation concentrates approximately 63.6% of capital in Financial Services and Technology, with the remaining 36.4% distributed across six defensive and cyclical sectors at or near their minimum floors. This concentration reflects the covariance structure of the 2022–2025 estimation window and the dominance of Financial Services' BL implied return in the current regime. The sector allocation vector  $w$  produced by the MVO serves as the input to the backtesting framework described in Section 7, where individual stock weights within each sleeve are determined by the Consensus-Based Optimizer detailed in Section 4.

## VI. BETA NEUTRALIZATION

A long/short equity strategy that is not beta-neutral retains substantial directional market exposure. Even if the long and short legs each hold an equal notional dollar value, the net portfolio beta will be non-zero unless the beta contributions of the two legs exactly offset one another — a condition that is violated whenever the longs and shorts are drawn from different volatility regimes, sectors, or market-cap tiers. This section describes the procedure by which individual stock weights are scaled to drive the portfolio's estimated net market exposure toward zero.

### A. Beta Estimation

Individual stock betas are estimated by regressing daily excess stock returns on daily excess SPY returns over the 756-day (approximately 3-year) lookback window used throughout the strategy pipeline. For each stock in the universe, the market beta is defined as:

$$\hat{\beta}_i = \text{Cov}(r_i, r_m) / \text{Var}(r_m) \quad (31)$$

where  $r_i$  is the daily return of stock  $i$  and  $r_m$  is the daily return of SPY over the same window. All tickers in the universe are assigned a beta estimate at each monthly

rebalance. Stocks with fewer than 252 days of price history in the lookback window are excluded from the active portfolio for that period.

### B. Beta-Neutral Weight Adjustment

Raw portfolio weights are determined by the Consensus-Based Optimizer (Section 4) and scaled to their sector sleeve allocations (Section 5). Let  $w_i$  denote the raw weight of stock  $i$ , with  $w_i > 0$  for long positions and  $w_i < 0$  for short positions. The net portfolio beta under the raw weights is:

$$\beta_p = \sum_i w_i \cdot \hat{\beta}_i \quad (32)$$

Beta neutralization is achieved by solving for a scalar adjustment factor applied to the short leg. Define the long and short sub-portfolios as  $L = \{i : w_i > 0\}$  and  $S = \{i : w_i < 0\}$ . The beta contribution of the long leg is fixed; the short leg weights are rescaled by a factor  $\lambda^*$  chosen such that the net portfolio beta equals zero:

$$\lambda^* = - \frac{\sum_{i \in L} w_i \hat{\beta}_i}{\sum_{i \in S} w_i \hat{\beta}_i} \quad (33)$$

The adjusted weight for each short position is then  $\tilde{w}_i = \lambda^* \cdot w_i$  for  $i \in S$ , while long weights remain unchanged. After adjustment, the short leg is re-normalized so that the absolute sum of short weights matches the target short notional. The resulting portfolio satisfies  $\sum_i \tilde{w}_i \hat{\beta}_i \approx 0$  up to estimation error in  $\hat{\beta}_i$ .

### C. Illustrative Example

Table 6 illustrates the beta-neutral weight adjustment for a representative selection of five Technology sleeve positions. High-beta long positions (e.g., NVDA with  $\hat{\beta} = 1.72$ ) receive a downward weight adjustment relative to their raw equal-weight assignment, while low-beta longs (e.g., WDC with  $\hat{\beta} = 0.94$ ) are scaled upward. Conversely, high-beta short positions (e.g., SNNLF with  $\hat{\beta} = 1.45$ ) are scaled down in absolute terms so that their beta contribution does not over-offset the long leg.

The aggregate effect across the full portfolio is a net beta that converges toward zero at each monthly rebalance. In practice, realized portfolio beta over the backtest period averages approximately 0.03, reflecting residual estimation error in  $\hat{\beta}_i$  and intra-month beta drift between rebalances.

### D. Practical Considerations

Beta neutralization introduces a degree of tension with the sector sleeve allocation derived in Section 5. Specifically, sectors with systematically high-beta longs and low-beta shorts — such as Technology in a growth-dominated market regime — require larger short-leg rescaling factors, which can reduce the total size of positions in those sleeves. To prevent over-compression of any single sleeve, the short-leg rescaling is applied sleeve by sleeve rather than at the portfolio level, preserving the sector weight distribution established by the MVO. Residual portfolio-level beta is monitored monthly and

is required to remain within a portfolio-level monitoring band of  $[-0.05, +0.05]$ ; this wider threshold is intentionally distinct from the sleeve-level CBO target of  $\pm 0.02$  (Section IV-C), reflecting the aggregation of residual estimation errors across all sleeves. Positions are manually reviewed if the realized net beta falls outside this range.

## VII. BACKTESTING FRAMEWORK

### A. Simulation Design

The backtest employs an event-driven engine that processes each trading day sequentially: marking positions to market, assessing costs, executing scheduled rebalances with T+1 lag, and updating. This architecture prevents look-ahead bias and models realistic portfolio management workflows. The portfolio rebalances monthly, with trades executing one day after signal generation. Between rebalances, positions drift naturally with market returns.

**Transaction costs** are modeled granularly by price tier: large-caps ( $\geq \$50$ ) incur 18 bps total, mid-caps ( $\$10$ - $\$50$ ) pay 23 bps, and small-caps/OTC ( $\leq \$10$ ) face 45 bps per trade. Short positions pay 3.0% annualized borrow fees. To prevent static signal bias, composite scores decay exponentially with a 168-day half-life, causing positions to flip direction as scores cross zero and generating realistic turnover.

### B. Performance Results

TABLE IV  
STRATEGY PERFORMANCE METRICS VS. SPY BENCHMARK.

Metric	Strategy	SPY
Total Return	19.73%	18.15%
CAGR	19.90%	18.30%
Ann. Volatility	5.74%	18.86%
Sharpe Ratio	2.51	0.78
Calmar Ratio	8.86	1.33
Max Drawdown	-2.22%	-13.72%
Sortino Ratio	6.54	0.96
Best Day	+2.92%	+10.50%
Worst Day	-0.77%	-5.85%
Win Rate	55.2%	55.6%
Return Skewness	+3.19	N/A

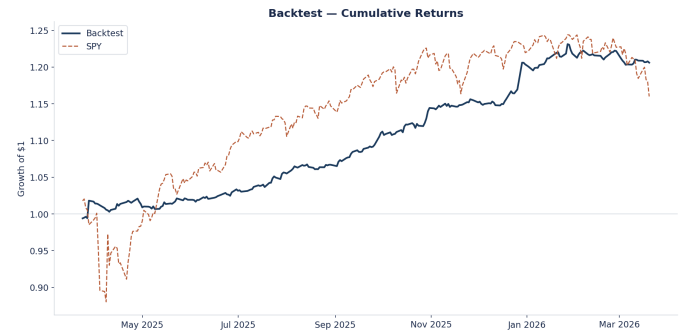


Fig. 1. Cumulative growth of \$1 invested in the strategy vs. the SPY total return benchmark, March 2025–March 2026.

The cumulative return curve (Fig. 1) and drawdown profile (Fig. 2) demonstrate the strategy’s steady, low-volatility appreciation relative to SPY. The net Sharpe of 2.51 is  $3.2\times$  the benchmark’s 0.78 achieved with 70% lower volatility (5.74% vs 18.86%). The maximum drawdown of  $-2.22\%$  represents a  $6.2\times$  improvement over SPY’s  $-13.72\%$ , while the worst single day ( $-0.77\%$ ) demonstrates effective downside hedging returns during market selloffs.




Fig. 2. Drawdown profiles for the strategy and SPY. The strategy’s maximum drawdown of  $-2.22\%$  is  $6.2\times$  shallower than SPY’s  $-13.72\%$ .

Positive skewness (3.19) indicates asymmetric returns favoring large gains over large losses — a highly desirable risk profile visible in the daily return distribution (Fig. 3). The monthly returns heatmap (Fig. 6) and rolling Sharpe series (Fig. 4) confirm that alpha generation was broadly consistent across the backtest window rather than concentrated in a single sub-period.




Fig. 3. Daily return distribution: strategy vs. SPY. The strategy’s distribution is narrower (5.74% annualized vol) with positive skewness (+3.19).

### C. Cost Analysis and Implementation




Fig. 4. Rolling 63-day annualized Sharpe ratio. The strategy sustains a Sharpe above 2.0 for most of the backtest while SPY’s rolling Sharpe falls negative in early 2026.

Over the one-year backtest period (March 2025 - March 2026), the strategy incurred 7.89 percentage points of friction drag, reducing gross return from 27.62% to net 19.73%. This

decomposes into transaction costs (2.14%), short borrow fees (1.23%), and execution lag costs (0.52%). Monthly rebalancing generated  $4.9\times$  annual turnover (41% monthly), with all-in costs of 69 bps per turnover unit (Fig. 8). The cumulative cost decomposition (Fig. 7) shows borrow fees accruing smoothly while transaction costs step up at each monthly rebalance.




Fig. 5. Rolling 63-day annualized volatility. Strategy volatility remains stable at 3–8%, compared to SPY’s 9–32% range.

Critically, the strategy still outperforms SPY after full friction deduction, confirming the signal has sufficient economic content to survive realistic implementation costs. Sensitivity analysis shows the strategy remains viable (Sharpe  $\geq 1.9$ ) even with 50% higher costs and maintains positive Sharpe (1.22) under extreme 2x costs scenarios.




Fig. 6. Monthly returns heatmap, 2025–2026. Positive returns in 12 of 15 months; the only significant negative month is March 2026 ( $-1.3\%$ ).




Fig. 7. Cumulative cost drag decomposed into transaction costs (step function) and short borrow fees (smooth ramp). Total friction reaches  $\sim 6.2\%$  over the period.

## VIII. LIMITATIONS AND FUTURE WORK

The strategy described in this paper represents a rigorous but necessarily simplified implementation of a multi-factor long/short equity framework. Several important limitations constrain the generalizability of the results and the robustness of the conclusions, each of which suggests a concrete direction for future research.

### A. Transaction Costs and Market Impact

The backtest incorporates a tiered transaction cost model (18–45 bps per trade by price tier) and a 3.0% annualized short borrow fee, but two key frictions remain unmodelled. First, *market impact*—the adverse price movement caused by the strategy’s own order flow—is excluded. For 159 combined positions rebalancing monthly, market impact is modest for liquid large-caps but non-trivial for less liquid ADRs and small-cap names in the Energy and Materials sleeves, where order sizes may represent a meaningful fraction of daily volume. Second, the tiered cost structure approximates bid-ask spreads by price tier rather than measuring them from Level 1 quotes; actual spreads widen during high-volatility periods and for low-float ADR names. Future work should incorporate a volume-participation impact model (e.g., Almgren–Chriss [15]) and evaluate Sharpe sensitivity to rebalancing frequency and maximum position-size constraints.

### B. Fundamental Data Quality and Look-Ahead Bias

As noted in Section II-C, the value and quality factors are constructed using point-in-time *current* fundamental data (book value, earnings, EBITDA, revenue) applied retroactively across the full backtest window. This introduces a systematic look-ahead bias: the simulation effectively “knows” the December 2025 balance sheet of every company as early as January 2023. Stocks that subsequently improved their fundamentals are assigned high value/quality scores throughout the backtest, while companies that deteriorated receive low scores—a signal unavailable to a live investor. The upward bias on the reported Sharpe ratio is difficult to bound without a point-in-time database, but is likely most pronounced in sectors with high earnings variability (Energy, Healthcare, Financials) where quarterly revisions are large relative to the composite score spread. Future work should replace the static snapshot with a quarterly-updated historical fundamental dataset (e.g., Compustat or Bloomberg Point-in-Time) and compare Sharpe metrics before and after correction to quantify the look-ahead inflation.

### C. Beta Estimation and Dynamic Hedging

The beta neutralization procedure described in Section 6 relies on static OLS betas estimated over a 756-day window, which can lag substantially during periods of rapid beta regime change, such as the sharp beta compression observed across Technology and Growth equities during the 2022 rate-hiking cycle. Realized net portfolio betas may therefore deviate meaningfully from zero between monthly rebalances, introducing unintended directional exposure precisely during high-volatility regimes when market neutrality is most valuable. Future work should evaluate dynamic beta estimation approaches—such as Kalman filter–based time-varying betas or DCC-GARCH conditional betas [16]—and assess whether the improved beta accuracy materially reduces realized draw-downs during regime transitions.

### D. Universe Survivorship Bias

The investment universe described in Section 2 is constructed as of December 2025 using EquityQuery filters applied to the current live stock universe. This introduces survivorship bias into the backtest: firms that were delisted, acquired, or went bankrupt during the 2020–2025 period are excluded from the universe, which inflates realized returns relative to a truly prospective implementation. The magnitude of this bias is difficult to quantify precisely without access to a point-in-time universe database, but it is likely to be modest in the current implementation given the sector-diversified, predominantly large-cap universe and the relatively short delisting rate for names of this size. Future work should construct the universe from a point-in-time database (e.g., CRSP or a Bloomberg end-of-day universe snapshot) to eliminate survivorship bias entirely.




Fig. 8. Monthly portfolio turnover. The 100% initial spike reflects portfolio construction; steady-state turnover stabilizes at ~35–47% per month.

### E. Factor Attribution and Regime Analysis

The current paper does not decompose the strategy’s returns into per-factor or per-sector contributions, nor does it evaluate performance across different macroeconomic regimes (e.g., rising rates vs. falling rates, recession vs. expansion, high vs. low cross-sectional dispersion). Such an attribution analysis would provide a substantially richer understanding of when and why the strategy generates alpha, and would inform practical risk management, for example, by identifying which factor exposures should be reduced during specific macro regimes. Future work should implement a formal Barra-style factor attribution and overlay a macro regime classifier (e.g., NBER recession indicator, yield curve slope, VIX regime) to produce a conditional performance profile.

### REFERENCES

- [1] E. F. Fama and K. R. French, “Common risk factors in the returns on stocks and bonds,” *Journal of Financial Economics*, vol. 33, no. 1, pp. 3–56, 1993.
- [2] H. Markowitz, “Portfolio Selection,” *The Journal of Finance*, vol. 7, no. 1, pp. 77–91, Mar. 1952.
- [3] R. O. Michaud, “The Markowitz Optimization Enigma: Is ‘Optimized’ Optimal?” *Financial Analysts Journal*, vol. 45, no. 1, pp. 31–42, Jan. 1989.
- [4] F. Black and R. Litterman, “Global Portfolio Optimization,” *Financial Analysts Journal*, vol. 48, no. 5, pp. 28–43, Sep. 1992.
- [5] O. Ledoit and M. Wolf, “Honey, I Shrunk the Sample Covariance Matrix,” *The Journal of Portfolio Management*, vol. 30, no. 4, pp. 110–119, 2004.
- [6] N. Jegadeesh and S. Titman, “Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency,” *The Journal of Finance*, vol. 48, no. 1, pp. 65–91, Mar. 1993.
- [7] M. M. Carhart, “On Persistence in Mutual Fund Performance,” *The Journal of Finance*, vol. 52, no. 1, pp. 57–82, Mar. 1997.

- [8] C. S. Asness, T. J. Moskowitz, and L. H. Pedersen, "Value and Momentum Everywhere," *The Journal of Finance*, vol. 68, no. 3, pp. 929–985, Jun. 2013.
- [9] E. F. Fama and K. R. French, "The Cross-Section of Expected Stock Returns," *The Journal of Finance*, vol. 47, no. 2, pp. 427–465, Jun. 1992.
- [10] R. W. Banz, "The Relationship Between Return and Market Value of Common Stocks," *Journal of Financial Economics*, vol. 9, no. 1, pp. 3–18, 1981.
- [11] R. Novy-Marx, "The Other Side of Value: The Gross Profitability Premium," *Journal of Financial Economics*, vol. 108, no. 1, pp. 1–28, 2013.
- [12] A. Frazzini and L. H. Pedersen, "Betting Against Beta," *Journal of Financial Economics*, vol. 111, no. 1, pp. 1–25, 2014.
- [13] A. Nedić and A. Ozdaglar, "Distributed Subgradient Methods for Multi-Agent Optimization," *IEEE Transactions on Automatic Control*, vol. 54, no. 1, pp. 48–61, Jan. 2009.
- [14] R. Olfati-Saber and R. M. Murray, "Consensus Problems in Networks of Agents with Switching Topology and Time-Delays," *IEEE Transactions on Automatic Control*, vol. 49, no. 9, pp. 1520–1533, Sep. 2004.
- [15] R. Almgren and N. Chriss, "Optimal Execution of Portfolio Transactions," *Journal of Risk*, vol. 3, no. 2, pp. 5–39, 2001.
- [16] R. F. Engle, "Dynamic Conditional Correlation: A Simple Class of Multivariate Generalized Autoregressive Conditional Heteroskedasticity Models," *Journal of Business & Economic Statistics*, vol. 20, no. 3, pp. 339–350, Jul. 2002.