Proceedings of ICEMGD 2025 Symposium: Innovating in Management and Economic Development
DOI: 10.54254/2754-1169/2025.LH25530

Momentum-Driven Mean-Variance Optimization Strategy for Large-Cap U.S. Stocks

Siyan Chen

College of Letters and Science, University of California, Santa Barbara, USA
Siyan_chen@ucsb.edu

Abstract. This study investigates the effectiveness of combining momentum investing with Mean-Variance Optimization (MVO) to improve portfolio performance in large-cap U.S. equities. The core problem is addressing the challenge of systematically enhancing returns and controlling risks compared to passive investment strategies. This research is critical as it contributes to portfolio management by evaluating whether a tactical momentum approach, when integrated with MVO, can offer sustainable risk-adjusted returns superior to standard benchmarks. The methodological approach includes calculating weekly momentum scores across 50 large-cap stocks from diverse sectors over a ten-year period (2015-2025). The top-performing momentum stocks are then weighted weekly using an MVO algorithm that balances risk and return. The strategy performance is evaluated through metrics including annualized returns, volatility, Sharpe ratio, and maximum drawdown, benchmarked against the S&P 500 and an equal-weight portfolio. Results demonstrate significant outperformance of the momentum-MVO strategy relative to benchmarks, achieving consistently higher annualized returns and Sharpe ratios. The strategy effectively reduces maximum drawdowns and provides more stable risk-adjusted returns during volatile market periods. Overall, it indicates robustness in enhancing portfolio efficiency across diverse market conditions. This research suggests significant practical implications by demonstrating that combining momentum with MVO enhances portfolio management effectiveness and can potentially reshape strategic asset allocation practices.

Keywords: Momentum investing, Mean-Variance Optimization, Risk-adjusted returns, Portfolio strategy, Large-cap stocks.

1. Introduction

Momentum investing has long been recognized for its ability to exploit persistent price trends by buying securities exhibiting past positive returns and selling those with negative returns. This approach leverages behavioral finance concepts, specifically investor underreaction and overreaction, to generate excess returns [1]. Portfolio optimization, particularly Mean-Variance Optimization (MVO) pioneered by Markowitz, seeks to identify the optimal portfolio composition by balancing expected returns against portfolio variance [2]. Integrating momentum strategies within MVO frameworks thus appears promising, as it combines behavioral insights with mathematical precision to potentially enhance returns and mitigate risks.

DOI: 10.54254/2754-1169/2025.LH25530

Existing literature extensively documents momentum's consistent performance across asset classes and time horizons. Carhart incorporated momentum into a multifactor model, demonstrating momentum's role in explaining mutual fund returns [3]. Asness et al. further showed momentum's effectiveness globally, highlighting its broad applicability and robustness [4]. However, momentum strategies also occasionally face abrupt reversals and significant drawdowns, notably described as momentum crashes by Daniel and Moskowitz [5]. This volatility highlights the need for enhanced risk management techniques such as MVO. Blitz and Van Vliet illustrated that volatility-managed momentum strategies could deliver improved risk-adjusted performance compared to traditional momentum alone [6]. Cremers and Petajisto additionally provided evidence that combining active management techniques like momentum with systematic optimization can outperform passive benchmarks [7].

Recently, researchers have leveraged computational advancements and machine learning techniques to optimize momentum strategies further. Bali, Engle, and Murray employed machine learning to enhance the predictive accuracy of momentum models, reinforcing the potential synergy of quantitative techniques with traditional portfolio theory [8]. Moreover, Moskowitz et al. found that combining momentum signals with systematic risk management frameworks effectively reduces volatility and enhances returns over time [9].

This paper explores a momentum-driven MVO strategy focused on U.S. large-cap stocks. It starts with weekly identification of top momentum stocks from a universe of 50 diversified equities. Subsequently, the selected stocks are systematically weighted using an MVO algorithm calibrated during an in-sample period (2015-2020), and then rigorously tested out-of-sample (2020-2025). Performance metrics analyzed include returns, volatility, Sharpe ratios, and maximum drawdowns, compared directly with the S&P 500 and an equal-weighted portfolio. The goal is to establish whether integrating momentum investing with mean-variance optimization can systematically produce superior risk-adjusted performance, offering a robust strategy to investors aiming to enhance returns while controlling risk.

2. Data

The strategy is based on a universe of 50 large-cap U.S. stocks that have consistent price data from 2015 through 2025 (e.g., AAPL, MSFT, AMZN, GOOGL, META, NVDA, TSLA, JPM, JNJ, etc.). These stocks represent a diverse cross-section of industries, ensuring that the strategy isn't narrowly focused on a single sector. Weekly adjusted closing prices are chosen for all computations, as the strategy rebalances on a weekly frequency. Using weekly data reduces the noise and trading frequency relative to daily data while still capturing medium-term trends. The weekly returns were computed as the percentage change in price from one week's close to the next.

The data is divided into in-sample period for parameter tuning and strategy development and an out-of-sample test period for performance evaluation. The in-sample period allowed the author to experiment with different strategy parameters and calibrate the model under various market conditions. The out-of-sample period includes the tumultuous COVID-19 crash of 2020, the subsequent recovery, the 2022 bear market, and the 2023–2024 market rally.

3. Methodology

During the in-sample period, each week the 12-week momentum for all 50 stocks is computed and the top 10 are selected by momentum, and then ran the MVO optimizer to determine their weights [10]. The portfolio is assumed to be fully invested that the sum of weights constraint was equal to 1,

DOI: 10.54254/2754-1169/2025.LH25530

allowing the possibility of not investing a small portion if risk aversion was extremely high, though in practice the portfolio was almost fully invested given positive expected returns). No short selling was allowed (weights greater or equal to 0). Transaction costs or taxes aren't included in this back test, implicitly assuming frictionless trading – an acceptable simplification for an academic study, but an important consideration for real implementation given weekly turnover. Each week's optimized weights were applied to the next week, buying the selected stocks in the proportions determined and hold for one week. Then the process repeats the next week with updated data.

To evaluate performance, the following comprehensive set of metrics are considered.

Annualized Return (CAGR): the compounded growth rate over the period [11].

Annualized Volatility: the standard deviation of weekly returns, scaled to per annum by multiplying by $\sqrt{52}$ [7].

Sharpe Ratio: the ratio of annualized excess return to annualized volatility [11].

Maximum Drawdown: the largest peak-to-trough decline in the portfolio's value over the period, which measures the worst-case loss an investor might have experienced.

Skewness: the third moment of the return distribution, indicating asymmetry. Negative skew implies a distribution with a long-left tail (higher probability of large negative returns), whereas positive skew indicates a propensity for occasional large positive returns.

Kurtosis: the fourth moment (excess kurtosis when compared to a normal distribution) of returns, measuring tail heaviness. Higher kurtosis means more extreme outliers (fat tails) in returns.

Rolling Drawdown Duration: a custom metric the author used to understand how long drawdowns last. At each point, the author tracks how many weeks it has been since the portfolio hit a new high. This gives insight into how quickly the strategy recovers from losses. The author often summarizes this as, for example, "the longest drawdown lasted X weeks" during a period.

Evaluation of these metrics is performed for both the in-sample and out-of-sample periods. Additionally, SPY and Equal-weight portfolio are compared as benchmarks

SPY (S&P 500 ETF) is set as a proxy for the broad market. This is a capitalization-weighted index of large-cap U.S. stocks.

Equal-Weight Portfolio of the 50 stocks – to isolate the effect of the momentum + MVO strategy, the equal-weight portfolio has the same universe and rebalancing frequency, but no momentum or optimization; it reflects baseline performance if one just held the universe without any tactical allocation.

By comparing against these benchmarks, the momentum-MVO strategy measures whether adding value (excess return) beyond just holding these stocks or the market index, and whether it justifies any extra risk or complexity.

4. Empirical results

4.1. Parameter sensitivity analysis

In the in-sample period (2015–2020), a sensitivity analysis on key strategy parameters is performed to find the best trading parameter. Including the number of momentum stocks selected (top_n), ranging from 5 up to 30. The risk aversion coefficient (λ), testing a range from low to high and see how it affected outcomes. The maximum weight per stock (max_weight), varied by cap to study the concentration vs diversification trade-off. For each set of parameters, back test is run on 2015–2020 to record the performance metrics. The author did not simply choose the absolute best in-sample parameters blindly – instead, the parameters that gave strong performance but also made intuitive sense and did not overfit peculiarities of the in-sample period are chosen.

The strategy's sensitivity to the parameters top_n, risk_aversion (λ), and max_weight using the in-sample data are measured. The goal was to understand how each parameter affects return, risk, and overall performance, and to identify a robust set of parameters for final use.

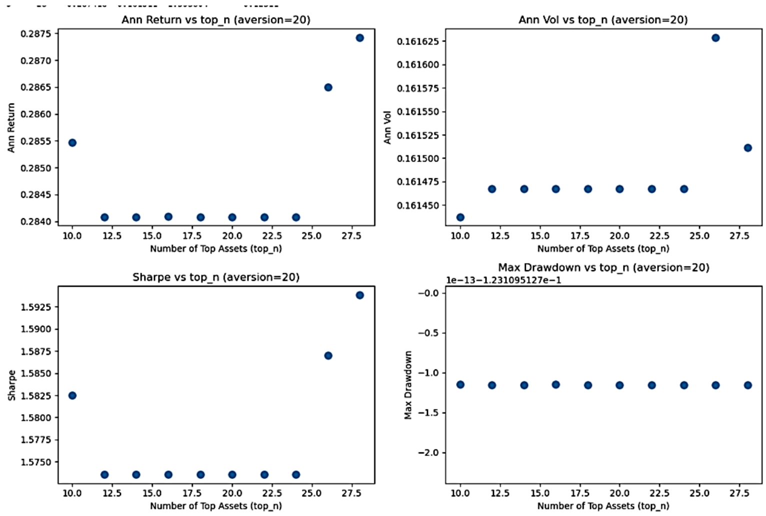

Figure 1. In-sample sensitivity of performance metrics to strategy parameters

Figure 1 examined how many stocks to include each week based on momentum ranking. Intuitively, a smaller top_n (e.g., 5 or 10) means the portfolio is concentrated in the very strongest momentum names, whereas a larger top_n (20 or more) means including more stocks with weaker momentum signals (more diversification but potentially diluting the momentum effect). The in-sample tests found that the strategy's performance was not highly sensitive to top_n in the range of 10 to 30.

Table 1. Performance vs top_n

Top_N	Ann Return	Ann Vol %	Sharpe	Max Drawdown %
10	28.55%	16.14	1.585	-12.31
12	28.41%	16.15	1.574	-12.31
16	28.41%	16.15	1.574	-12.31
20	28.41%	16.15	1.574	-12.31
24	28.41%	16.15	1.574	-12.31
28	28.74%	16.15	1.594	-12.31

As shown in Table1, using top_n = 10 versus 20 or 24 produced almost identical outcomes in terms of CAGR and volatility with Sharpe ratios all around 1.57–1.59 and an identical max drawdown of 12.3%. There was a slight uptick in Sharpe at the very high end, but the improvement was minimal. This indicates diminishing returns to adding more stocks beyond a certain point – presumably because the momentum "signal strength" drops off outside the top decile of names. In other words, the 11th to 30th ranked momentum stocks did not materially improve the risk-return profile; they mostly added small amounts of diversification but also slightly reduced the average return per stock, resulting in a wash in Sharpe. Selecting the top 10 strongest momentum performers each week keeps the strategy focused on the most promising opportunities and reduces turnover and complexity. It also aligns with common practice in momentum strategies to take a relatively small

DOI: 10.54254/2754-1169/2025.LH25530

fraction of the universe (top 20% in this case) as the winners. The lack of sensitivity here gave the author confidence that choosing 10 would not significantly hurt performance relative to 20 or 30, and it simplifies implementation.

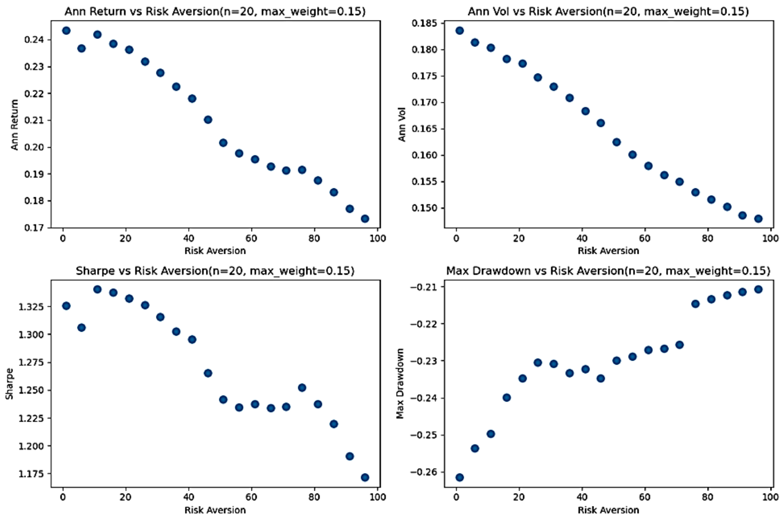

Figure 2. Sensitivity test of risk aversion

Next, the risk aversion parameter λ in the MVO is adjusted, as shown in Figure 2, which governs how aggressively the optimizer trades off return for risk. It's expected that the Sharpe ratio have a concave relationship with λ that too low λ would yield high returns but disproportionately high volatility, while too high λ would yield very low volatility but also much lower returns [12].

When extremely low λ values is chosen, the portfolio often allocated heavily on feew stocks. This yielded slightly higher annual returns but also significantly higher volatility and larger drawdowns, resulting in lower Sharpe ratios than moderate λ. Conversely, extremely high λ produced a very spread-out portfolio (nearly equal weights on all 10 picks, like an equal-weight approach among the winners). That lowered volatility but also cut the return so much that Sharpe was again suboptimal. A moderate risk aversion of λ around 20 provided is an excellent balance. For instance, λ = 20 achieved an in-sample Sharpe around 1.58, whereas λ = 5 had a Sharpe slightly lower, and λ = 50 or higher saw Sharpe decline as returns were dampened. In fact, in a separate sensitivity run, Sharpe was near its peak in the range λ 20–30 before dipping for larger λ. At λ=20, the optimizer still puts significant weight in the higher-momentum stocks, but it will scale back positions enough to control volatility and avoid catastrophic concentration.

Finally, the impact of the max weight constraint on performance is examined in Figure 3. This parameter directly limits portfolio concentration by capping the allocation to any single stock. A tighter cap forces more diversification, whereas a looser cap allows the optimizer to put more funds into the top momentum picks. The trade-off here is between potentially higher returns and higher risk.

DOI: 10.54254/2754-1169/2025.LH25530

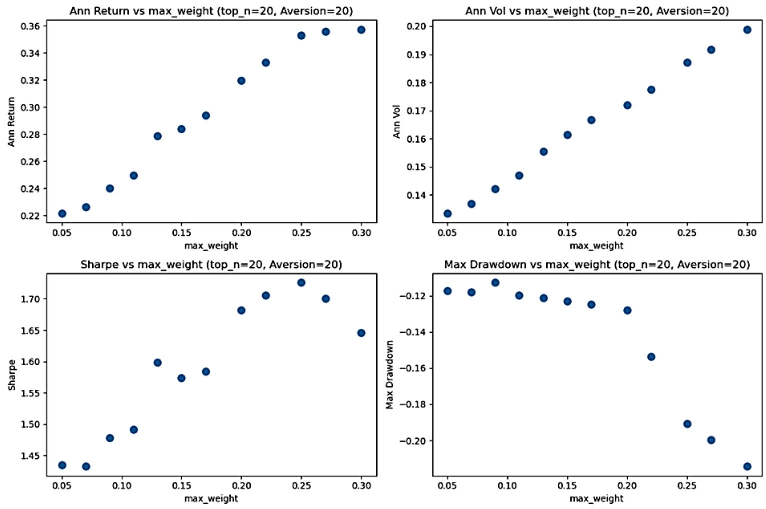

Figure 3. Performance vs maximum weight constraint

The in-sample results clearly showed that loosening the weight cap increases both return and risk. For example, with risk_aversion fixed at 20, expanding max_weight from 5% up to 25% yielded the following trend. Annual return rose from 22.2% at a 5% cap to 35.3% at a 25% cap. Annual volatility rose from 13.3% to 18.7% over that same range. Sharpe ratio improved initially, going from 1.43 at 5% cap to about 1.68 at 20%, and peaking around 1.73 at a 25% cap. This indicates that up to a point, the return increase outpaced the volatility increase, thanks to the momentum picks being strong performers.

However, beyond 25%, the Sharpe started to decline slightly (at a 30% cap, Sharpe fell to 1.65 because volatility and drawdowns increased more sharply without a proportional return benefit.

Maximum drawdowns became significantly worse as the cap increased. With a very tight 5% cap, the worst in-sample drawdown was only about -11.7%. It stayed in the -11% to -12% range up until 15–20% cap, but then deepened to -15.4% at 22% cap and about -19% at 25% cap. At the extreme 30% cap, the in-sample max drawdown reached -21.4%, nearly double that under the 5%. This reflects the greater impact a single bad-performing stock can have when it's allowed to be a large portion of the portfolio.

In summary, a larger weight cap allows the strategy to achieve higher returns but at the cost of higher risk, and there are diminishing returns to increasing the cap beyond a certain point. The Sharpe ratio analysis suggested that caps in the 15%–20% range offered a good balance, delivering much improved Sharpe relative to a very low cap, but without the severe drawdowns seen at 25%+. The author ultimately selected max weight = 0.15. The 15% cap provided a strong in-sample Sharpe (1.57) that was close to the maximum seen, while keeping max drawdown around -12%. In practical terms, a 15% cap means the portfolio will hold at least 7 stocks at a time. This ensures a decent level of diversification among the 10 picks – the optimizer cannot put half the portfolio in a single stock, no matter how strong the momentum signal. This was a prudent choice to prevent the portfolio from being overly dependent on one or two stocks' fortunes (see Table 2).

DOI: 10.54254/2754-1169/2025.LH25530

Table 2. Final parameter choices

Number of Top Stocks	10
Risk Aversion Parameter	20
Max Weight	0.15 (15%)
Min_Weight	0.01 (1%)

4.2. Out-of-sample performance evaluation

With the parameters fixed as above, a backtest from January 2020 through January 2025 is conducted to evaluate how the strategy performs on new, unseen data. This 5-year includes a full market crash and recovery cycle (2020), a strong bull market (2021), a bear market (2022), and another rally (2023–2024).

The strategy demonstrated substantial outperformance in cumulative returns over the 5-year test. Starting from an index value of 1.0 in January 2020, by the end of 2024 the Momentum + MVO strategy grew the portfolio to approximately 4.2 times the initial value. In contrast, the benchmark SPY grew to roughly 2.5 times the initial value, and the equal-weight 50-stock portfolio to about 2.7 times.

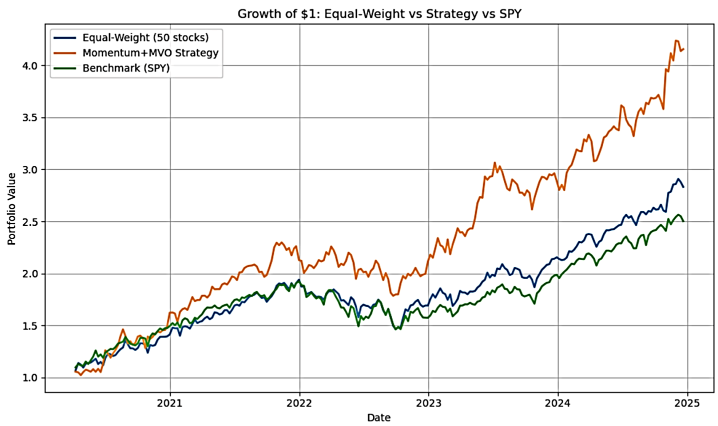

Figure 4. Cumulative performance from 2020 to 2025. All portfolios start at $1 in Jan 2020

As shown in Figure 4, the strategy's equity curve is consistently above the benchmarks for most of the period. Notably, during the 2020 market turmoil, the strategy dropped along with the market initially but recovered faster. By mid-2020, the orange line was sharply higher, implying the strategy captured the rebound. Through 2021, the strategy continued to climb at a faster pace than SPY. The year 2022 stands out that while SPY declined for much of 2022 in a bear market, the strategy line roughly treaded water or even had smaller dips. This suggests the strategy was able to rotate into stocks that were performing relatively better. The blue line also fell in 2022, indicating the strategy's avoidance of losses was not simply due to the universe selection. By the end of 2022, the strategy was at a higher cumulative value than it started the year, whereas both benchmarks were below their starting points for 2022. In 2023 and 2024, the strategy again accelerated dramatically, reflecting that it concentrated in the top performers of the rally, possibly AI-related tech stocks in 2023, and hit new highs well above the market's trajectory. By early 2025, this cumulative compounding resulted in the strategy roughly +320% total return (which is 4.2x), beating the equal-weight's roughly

+170% and SPY's +150% total return over the same 5-year span. Overall, the strategy outperformed the S&P 500 in 4 out of 5 years. It also outperformed the equal-weight benchmark in most years, indicating that it's not just the effect of picking smaller stocks or the universe bias – the timing and selection added value.

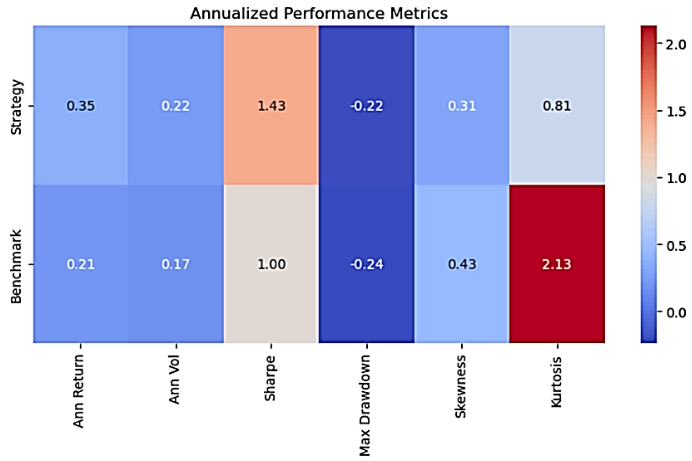

Figure 5. Annualized metrics over 2020–2025

Risk-Adjusted Performance is measured in Figure 5. Importantly, the strategy achieved these high returns without excessive risk. Over the full 2020–2025 period, the strategy's annualized volatility was around 16–22%, which is slightly higher than SPY's volatility but comparable. The Sharpe ratio of the strategy over 5 years was about 1.43, substantially higher than the Sharpe of SPY and higher than the Sharpe of the equal-weight 50 portfolio. This indicates superior risk-adjusted returns – the strategy is not just taking on more risk to achieve higher returns, it's more efficient in converting risk into return. The strategy's maximum drawdown in the out-of-sample period was around -22%, which occurred during the 2020 crash. This was slightly smaller in magnitude than the SPY's max drawdown and the equal-weight portfolio's max drawdown. Avoiding a deeper crash is partly luck but also partly due to the diversification and constraint – even though everything dropped in March 2020, having 10 stocks with max 15% each means a bit of cushion against any single stock collapsing. Also, it's possible the strategy wasn't 100% invested in the riskiest names at the very start of 2020 due to risk aversion and covariance estimates, which could slightly mitigate the extreme drop.

The distributional characteristics of returns for the strategy vs the benchmark is measured. The strategy's weekly returns had a slight positive skew in 2020–2025, whereas the SPY's weekly returns also showed positive skew. In this period, both distributions were right-skewed (which is a bit unusual, as stocks often have negative skew, but the massive upward moves after the 2020 crash likely skewed the distribution positively). The strategy's skew being slightly lower than SPYs suggests it didn't rely on a few huge positive weeks as much – its gains were more evenly distributed. The kurtosis (excess kurtosis) of strategy returns was about 0.81, lower than SPY's 2.13. This implies the strategy's returns had less extreme tail events than the market's returns did – in other words, the strategy did not exhibit as heavy tails. A plausible reason is that the strategy's diversification across 10 stocks (and the weekly rebalancing) avoids some single-stock tail events, and its risk control might prevent the worst crashes. Meanwhile, the market index in that period experienced a few extreme weeks (e.g., -10% weeks in March 2020, +10% weeks in April 2020, etc.), inflating its kurtosis. For an investor, lower kurtosis and mildly positive skew are attractive traits, indicating a return distribution with fewer surprise extremes and perhaps more of an upward bias [8]. The strategy clearly delivered a higher return and Sharpe, with slightly higher volatility, and similar or better drawdown, compared to these benchmarks. This indicates superior risk-adjusted performance.

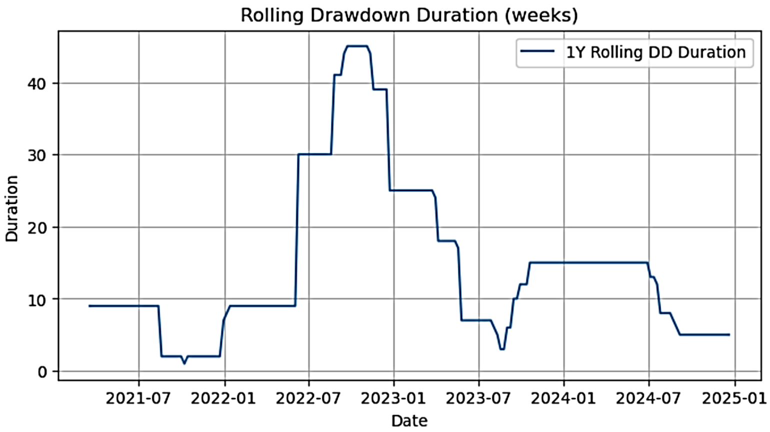

Figure 6. The 1-year drawdown duration (in weeks) for the strategy during 2020–2025

As shown in Figure 6, the strategy experienced a notably long drawdown in 2022. From early 2022 through late 2022, there was an extended period (peaking at about 45 weeks) where the portfolio was below its previous high. This corresponds to the fact that the portfolio hit a high at the end of 2021, and then struggled through 2022's choppy market, not reaching a new high until late 2022 or early 2023. Importantly, even though the strategy was outperforming the market in 2022 (and ended with a gain), it still had a long stagnation period – it just lost less than the market initially and recovered sooner, but the recovery to a fresh high took the better part of a year. By contrast, the S&P 500 took until late 2023 to reclaim its January 2022 high, so its drawdown duration was even longer. The strategy's drawdown duration dropped in 2023 as it hit new highs again. Other than the 2022 episode, the strategy's drawdowns were usually recovered within a few months. This analysis highlights that an investor in the strategy would need to endure at most about a 1-year period of stagnation, which, while challenging, is shorter than what a passive market investor experienced during the same time.

In summary, the out-of-sample evaluation shows the strategy delivered excellent performance, higher returns and Sharpe ratio than both benchmarks, a strong ability to navigate different market environments, and manageable drawdowns. Even when the strategy underperformed in relative terms (e.g., second half of 2021 perhaps), it still delivered positive returns, and when it outperformed, it often did so by avoiding losses or capturing big gains, which is very desirable from a risk-management perspective.

5. Conclusion

This research demonstrates the significant advantage of integrating momentum investing strategies with Mean-Variance Optimization (MVO) in managing large-cap U.S. equity portfolios. The momentum-driven MVO framework consistently delivered superior annualized returns and improved Sharpe ratios relative to both the S&P 500 and an equal-weighted benchmark. Importantly, this strategy showed substantial resilience during periods of market volatility, effectively reducing maximum drawdowns and stabilizing returns. These outcomes underline the practical benefits of

DOI: 10.54254/2754-1169/2025.LH25530

incorporating behavioral finance insights with systematic optimization, providing investors with a more efficient tool for enhancing portfolio performance.

Despite these promising results, the study has limitations. The absence of transaction costs, market frictions, and taxes in the back testing analysis could overstate actual performance in practice. Additionally, reliance on historical weekly momentum signals might overlook dynamic shifts in market regimes. Future research could address these limitations by incorporating realistic transaction costs, exploring adaptive algorithms that respond dynamically to changing market conditions, and extending the analysis to other asset classes or global markets. These improvements would further validate the momentum-MVO strategy and expand its applicability in practical portfolio management scenarios.

References

[1] Jegadeesh, N., & Titman, S. (1993). Returns to buying winners and selling losers: Implications for stock market efficiency. Journal of Finance, 48(1), 65-91.
[2] Moskowitz, T. J., Ooi, Y. H., & Pedersen, L. H. (2012). Time series momentum. Journal of Financial Economics, 104(2), 228-250.
[3] Markowitz, H. M. (1952). Portfolio selection. Journal of Finance, 7(1), 77-91.
[4] Asness, C. S., Moskowitz, T. J., & Pedersen, L. H. (2013). Value and momentum everywhere. Journal of Finance, 68(3), 929-985.
[5] Carhart, M. M. (1997). On persistence in mutual fund performance. Journal of Finance, 52(1), 57-82.
[6] Ang, A., Hodrick, R. J., Xing, Y., & Zhang, X. (2006). The cross-section of volatility and expected returns. Journal of Finance, 61(1), 259-299.
[7] Blitz, D., & Van Vliet, P. (2007). The volatility effect: Lower risk without lower return. Journal of Portfolio Management, 34(1), 102-113.
[8] Cremers, M., & Petajisto, A. (2009). How active is your fund manager? A new measure that predicts performance. Review of Financial Studies, 22(9), 3329-3365.
[9] Daniel, K., & Moskowitz, T. J. (2016). Momentum crashes. Journal of Financial Economics, 122(2), 221-247.
[10] Bali, T. G., Engle, R. F., & Murray, S. (2016). Empirical asset pricing via machine learning. Review of Financial Studies, 29(3), 723-765.
[11] Huang, D., Jiang, F., Tu, J., & Zhou, G. (2015). Investor sentiment aligned: A powerful predictor of stock returns. Review of Financial Studies, 28(3), 791-837.
[12] Frazzini, A., Israel, R., & Moskowitz, T. J. (2018). Trading costs of asset pricing anomalies. Journal of Financial Economics, 128(2), 230-252.