Volatility Model Selection and Momentum Performance

Implications for Dynamic Strategies

Adam Gyalogh* Ante Jälevik**

Master's Thesis in Finance
Stockholm School of Economics
May 2025

Abstract

This thesis investigates the role of volatility forecasting in enhancing momentum strategy performance through a dynamic weighting framework. Building on the work of Daniel and Moskowitz (2016), the study evaluates whether more sophisticated volatility models lead to meaningful improvements in risk-adjusted returns. Six volatility forecasting methods, including realized variance and five GARCH-type models', are applied within a consistent dynamic momentum strategy across U.S. and U.K. equity markets. Portfolio weights are constrained to lie between 0 and 1, and the lambda parameter controlling volatility exposure is estimated ex-ante using a rolling window. While dynamic strategies generally improve Sharpe ratios compared to a static momentum benchmark, the results show that higher forecasting accuracy does not consistently yield higher returns. Notably, simpler models such as realized variance perform competitively, raising questions about the marginal value of model complexity. The findings emphasize the trade-off between forecast precision and practical implementation, suggesting that model selection may be best guided by investor risk preferences rather than predictive accuracy alone.

Keywords: Momentum factor, Dynamic strategy, Volatility forecast

Supervisor: Magnus Dahlquist, Professor of Finance, Department of Finance, Stockholm School of Economics

Acknowledgements: We would like to express our gratitude to our supervisor, Magnus Dahlquist, for his guidance and support throughout the development of this thesis.

* Student number: 42686
** Student number: 42639

Table of Contents

1 Introduction 1

2 Literature Review 5

3 Data Description 8

4 Methodology 10

4.1 Models 12

5 Analysis of the Momentum Performance 15

6 Results 21

6.1 Volatility forecasts 24

6.2 Dynamic strategy performance 25

7 Discussion 34

8 Conclusion 36

9 References 37

9.1 Papers 37

9.2 Other 38

10 Appendix 39

10.1 Discussion of the Behavior of Lambda 39

10.2 Trading Cost Considerations 42

10.3 Six-month momentum portfolio results 46

10.4 Price Series Cleaning Procedure for Delisted Securities 47

10.5 Artificial Intelligence Disclosure 47

1 Introduction

Since it was first introduced by Jegadeesh and Titman in 1993, the momentum factor has remained a subject of debate in academic finance, particularly regarding how to explain the persistence in stock returns. The original explanatory argument of markets underreacting to positive or negative information surrounding a stock has been expanded on in countless ways. More than thirty years later, researchers still have not reached consensus on a single, dominant framework that fully explains momentum-driven returns.

In parallel with theoretical debates, there is an abundance of approaches to implementing momentum as an investment strategy. These strategies differ based on portfolio formation rules, stock universe selection, holding period, and the type of information incorporated. This paper examines the role of volatility forecasts in enhancing the performance of a dynamically weighted momentum strategy.

The discourse surrounding momentum can be categorized in several ways. Some explanations are rooted in behavioral finance (eg., Barberis, Shleifer, and Vishny, 1998), while others draw from more traditional, rational frameworks. Certain scholars argue that the momentum effect operates primarily at the factor or industry level, rather than at the individual stock level (Gupta & Kelly, 2018; Scowcroft & Sefton, 2005). Others contend that momentum returns represent compensation for bearing additional risk (eg., Kelly, Moskowitz and Pruitt, 2021). Through certain periods, momentum underperformed relative to the market, prompting skepticism about the persistence of the factor. Yet, over time, momentum has again proved that it has not lost what it was claimed to be.

One key characteristic of the momentum factor is its sensitivity to volatile markets. Daniel and Moskowitz (2016) illustrate how the momentum strategy has historically suffered significant losses following market crashes. They show how a dynamic momentum strategy can help mitigate this weakness and improve the Sharpe ratio compared to a static strategy. In contrast to Daniel and Moskowitz's complex weighting model, Barroso and Santa-Clara (2015) propose a much simpler constant volatility approach, which also outperforms the traditional static momentum strategy. These examples highlight how volatility can be leveraged to increase risk-adjusted returns for the momentum strategy.

This paper extends their analysis by exploring how different approaches to volatility modelling affect the performance of a dynamic momentum strategy.

Striking a balance between theoretical improvements and practical implementability is a delicate challenge. Complex volatility models often depend on strong assumptions about the structure of data. Further, more components in a dynamic weighting system introduces more risk for biased estimates. In a theoretical setting these parameters and inputs can be tweaked and adjusted to produce favorable results, in a practical setting however, backtesting has inherent limitations.

This paper discusses the trade-off between complexity and performance in volatility forecasting. Specifically, it investigates the extent to which more sophisticated volatility models improve forecast precision, and whether this added precision translates into meaningful gains in risk-adjusted returns.

We begin the analysis by implementing a standard momentum strategy based on the methodology of Jegadeesh and Titman (1993), which serves as a baseline for evaluating enhancements through dynamic weighting. Building on this, a dynamic momentum strategy is constructed following the approach of Daniel and Moskowitz (2016), in which the portfolio allocation between the momentum portfolio and a risk-free asset is determined by a dynamic weighting function using forecasts of expected returns and conditional volatility. These strategies are applied to two distinct geographical stock universes, one comprising of stocks listed on the New York Stock Exchange and Nasdaq (US) and the other of stocks listed on the London Stock Exchange (LSE). Comparing the two markets allows for an assessment of the robustness and generalizability of our findings across different market environments. Our implementation differs from these previous studies in that returns are based solely on price data and do not account for dividends or distributions. While this limits direct comparability with total-return-based studies, it does not affect the internal validity of our analysis, which focuses on the relative performance of volatility models.

The main contribution of this paper is to evaluate the effect of volatility model selection on risk-adjusted returns of momentum investing. Compared to the approach of Daniel and Moskowitz (2016), we make a number of additional adjustments to the implementation of the dynamic weighting function. Firstly, we modify the method used to determine the parameter lambda, which governs the scaling of portfolio volatility within the dynamic weighting function. In order to adhere to a strict ex-ante approach, we estimate this parameter on a rolling monthly basis using data for the preceding five years rather than setting a constant value throughout, based on information from the entire observed period. Secondly, we constrain the

portfolio weights to lie within the range [0, 1], to better reflect the practical limitations typically faced by real-world investors.

Our momentum portfolio volatility forecasts are generated using six different models; one realized variance model and five different specifications defined to better capture certain characteristics of the data of generalized autoregressive conditional heteroskedasticity (GARCH) models. These models differ in complexity and modelling assumptions. Some are designed to capture features such as asymmetry or skewness in volatility, while others aim to incorporate exogenous information for improved forecasting accuracy.

By applying these models within a consistent momentum framework, the analysis isolates the effect of volatility forecast quality on portfolio performance, shedding light on whether model selection results in tangible improvements for risk-adjusted returns.

We find that all six evaluated volatility models, when applied within the dynamic strategy framework, improve the Sharpe ratio compared to the standard (static) momentum strategy when applied to our US stock universe. For the London stock universe four out of six models produce the same improvements. The most efficient model in terms of risk-adjusted returns is the EGARCH for the US markets, and basic GARCH (1,1) for the LSE. Although improvements are observed, the differences in Sharpe ratio between the models are marginal.

Furthermore, we find that models with the highest volatility forecast accuracy do not necessarily yield the best risk-adjusted returns. Our findings suggest that certain model characteristics, such as the high persistence of volatility forecasts in the realized variance model, can enhance strategy performance under certain market conditions. This implies that forecasting accuracy alone is not a sufficient criterion for model selection in a dynamic momentum framework.

In addition, we find that the dynamic weighting function works as intended. As we scale the lambda to match the dynamic momentum volatility to market volatility and constrain weights on the momentum portfolio to be between 0 and 1, we observe an interesting pattern in the resulting allocations between the momentum portfolio and risk-free asset. When the market portfolio volatility exceeds the standard momentum portfolio volatility, the weighting function takes on an almost exclusively binary form, shifting between a full allocation to either asset depending on expected returns and forecasted volatility. When instead the standard momentum portfolio volatility exceeds the market portfolio volatility, the weights of the dynamic strategy

start to take on a more nuanced range of values, signaling that our ex-ante estimation of the lambda parameter in the optimal weight function works well.

Our results indicate that model selection for volatility forecasting does not provide any large gains in risk-adjusted returns and a simple volatility model can be sufficient. Perhaps the most important finding of this paper is the demonstration of how well momentum volatility can be managed using this dynamic weighting framework. Our application of this strategy leads to significant reductions in volatility compared to the static momentum strategy. Further, the lambda parameter can be utilized to determine the approximate level of volatility for the dynamic momentum portfolio, allowing for better control of downside risk typically faced by momentum investors.

2 Literature Review

The momentum factor was first formally introduced by Jegadeesh and Titman in 1993. Their paper demonstrates that stocks that have performed well over the past three to twelve months tend to continue to outperform the market in the subsequent period, while past losers continue to underperform. This finding challenges the notion of weak-form efficiency in financial markets, suggesting that past price trends may hold predictive power for future returns. Jegadeesh and Titman evaluate varying time frames, constructing 16 strategies in total with lookback and holding periods of 1, 2, 3 and 4 quarters. Further, they construct another 16 strategies where they skip a week between the lookback and formation period. By constructing zero-cost portfolios that go long on past winners and short on past losers, the authors provide robust empirical evidence of the momentum effect. The paper laid the foundation for extensive subsequent research on momentum's persistence, risk-based and behavioral explanations.

A few years after the academic introduction of the momentum factor, Barberis, Shleifer and Vishny (1997) examine why the observed persistence, or momentum, in stock returns exists. They also analyze a reverse mechanism, in which over a longer horizon of three to five years, security prices tend to overreact to consistent news in the same direction, whether positive or negative. To explain these irregularities, they look at how investors form their beliefs, relating this to classical behavioral heuristics such as representativeness and conservatism. They especially connect conservatism to the momentum factor, a phenomenon where agents have been shown to be slow in adapting to new evidence. This represents one of the early attempts to explain momentum through the lens of behavioral finance.

While Barberis, Shleifer, and Vishny (1997) offer a theoretical explanation rooted in investor psychology, Cooper, Gutierrez, and Hameed (2004) approach the same areas from a more empirical perspective, focusing on how momentum behaves across different market states. By defining market periods as positive or negative based on returns, they show that the mean monthly momentum profit differs based on market state. The two market states are defined based on the lagged three-year market return, classified as "UP" if the return is non-negative, and "DOWN" if it is negative. Over longer horizons, they find that momentum profits, whether following up or down markets, tend to reverse, consistent with long-term return overreaction.

The following research on momentum has taken many different angles in its approach to understanding the phenomenon. Scowcroft and Sefton (2005) show that momentum exists

outside the idiosyncratic stock-level. Scowcroft and Sefton identify a momentum phenomenon on an industry level. This contrasts with the previously established view of momentum as a characteristic of individual stocks. By composing winner and loser portfolios based on industry sorting, the authors show that in a value-weighted large-cap universe momentum returns are mainly driven by industry momentum. In contrast, the momentum returns in a small-cap universe are shown to primarily be driven by individual stock momentum.

Later, Gupta and Kelly (2018) would similarly find that momentum exists on a factor level. By analyzing 65 characteristic-based equity factors, they show that individual factors experience persistence in returns. This persistence can be used to time factor focused investments and the authors show that this phenomenon is a key driver of momentum in general. The paper builds on these findings by constructing an alternative momentum-based investment strategy which combines traditional momentum investing with an approach that considers factor momentum and the value factor. This strategy is shown to outperform a momentum strategy that solely focuses on stock-level momentum, using the Sharpe ratio as the evaluation metric.

Kelly, Moskowitz and Pruitt (2021) evaluate momentum factor returns from a risk-based perspective. They explore a new approach to explain momentum returns that captures previously missed risk premia associated with a momentum strategy. A principal component analysis is applied to show that loadings on other factors vary significantly even over short time periods for a momentum portfolio. The authors conclude that this exposure to conditional risk can explain a large fraction of previously theoretically unmotivated momentum returns.

All these different approaches to understanding momentum have laid the groundwork for various methods of maximizing risk-adjusted momentum returns, Daniel and Moskowitz (2016) present one such strategy. The authors show that by applying a dynamic portfolio weighting the effectiveness, measured by the Sharpe ratio, of the momentum portfolio can be increased. The paper starts by establishing a few key characteristics of the momentum factor. Importantly, Daniel and Moskowitz analyze the poor performance of the momentum strategy following market crashes. Specifically, they show that the rebound typically experienced by losing stocks after a crash is a key factor explaining this underperformance. To mitigate the negative impact of market turmoil on momentum returns, the authors use forecasts of momentum volatility. They determine the portfolio assortment between the risk-free asset and the winner minus loser portfolio so that the dynamic strategy's conditional volatility is proportional to the conditional Sharpe ratio. A generalized autoregressive conditional

heteroscedasticity (GARCH) model is used to forecast volatility, which is then in turn used for dynamic weighting.

Another method of utilizing volatility to increase profits from a momentum strategy is presented by Barroso and Santa-Clara (2015). They find that the risk of momentum is predictable and varies over time. The authors apply a constant volatility strategy which is shown to nearly double the Sharpe ratio of the momentum strategy. This result is achieved by scaling their strategy's risk to be constant over time, using a realized variance-based forecast of variance. Barroso and Santa-Clara's findings, along with those of Daniel and Moskowitz, illustrate the benefits of incorporating methods to utilize the predictability of momentum volatility.

Moreira and Muir (2016) also explore the topic of utilizing volatility to increase the effectiveness of a portfolio. The authors show that volatility timing increases the Sharpe ratio for, among others, the momentum factor. According to Moreira and Muir this increased effectiveness can be explained by the fact that changes in volatility are not offset proportionally by changes in expected returns. These results are achieved by decreasing the risk exposure of the portfolio when variance has been high in the previous period. In contrast to the strategy of Daniel and Moskowitz (2016) this approach does not require forecasting of the portfolio volatility.

3 Data Description

Our data is sourced from the Refinitiv Eikon Database, accessed via the SSE Campus computer during February and March of 2025. We have collected every International Securities Identification Number (ISIN) from the New York Stock Exchange (NYSE) - United States, Nasdaq - United States, and the London Stock Exchange (LSE) - United Kingdom, covering all ISINs that have ever existed. Using these ISIN codes, we retrieved the stocks' price series between 1st of January 1973 and 31st of December 2024. The approximate number of unique securities is 12 900 for Nasdaq and NYSE, and 4 500 for the London Stock Exchange.

All daily returns are calculated as close-to-close returns, based solely on price data and do not include dividends or other distributions. As a result, our return measures reflect capital gains only, which may limit direct comparability with studies that use total return data (e.g., Daniel and Moskowitz, 2016). The stock price series are adjusted for stock splits, reverse splits, and bonus issues, providing a split-adjusted basis that ensures consistency in historical price comparisons.

Not every security will have a price for the entire date range. However, the number of securities available each day is sufficiently large to allow for meaningful analysis.

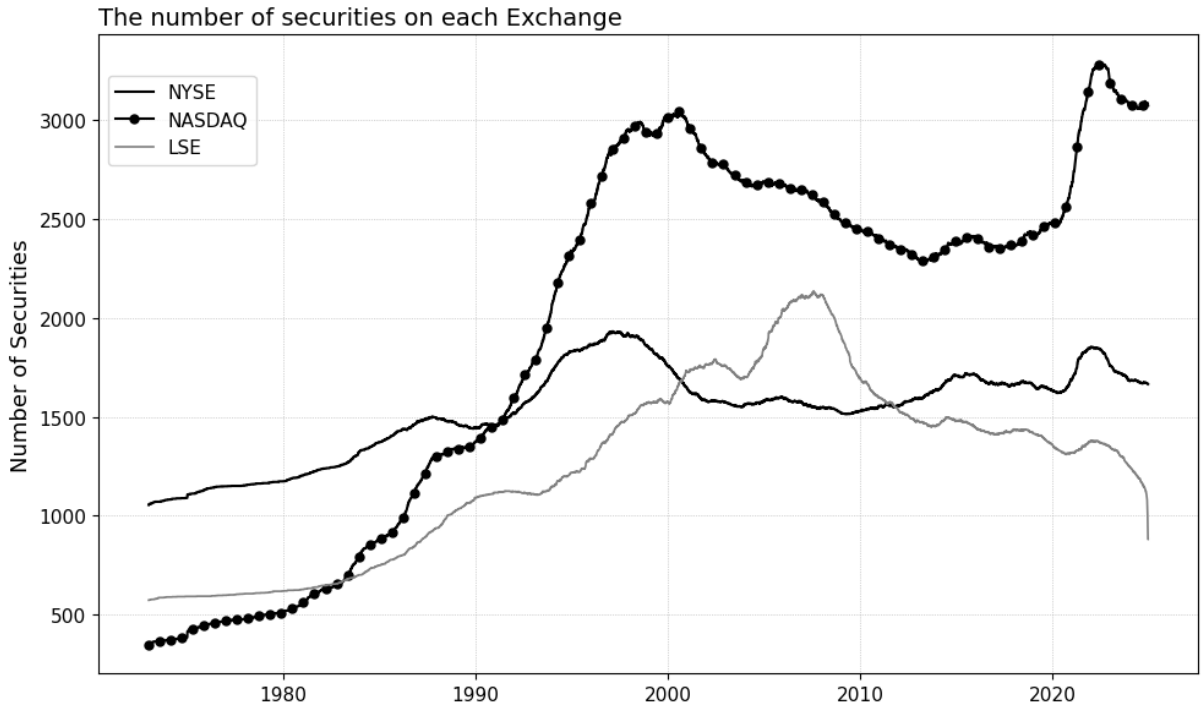


Figure 1. The number of securities listed across time on each exchange, based on data sourced from Refinitiv Eikon

Refinitiv Eikon also includes a special solution for securities that are no longer listed on exchanges, such as those that have been delisted or merged. For these securities, the last known price is repeated daily until the end of the price series. We identify and adjust these cases to avoid distortions in our analysis of security returns when sorting stocks into portfolios. Specifically, we exclude these artificially extended portions of the series from our dataset, as the repeated prices yield 0% returns that can mislead the sorting algorithm. For instance, during a market downturn, a security showing a 0% return might incorrectly appear to be a top performer, resulting in its inclusion in the "winners" group. Conversely, in a bull market, the same security could be misclassified as a "loser." In both cases, such a selection would be incorrect, since the 0% return reflects a database artifact rather than actual market performance.

To ensure reproducibility, the code used to process the data is provided in the appendix. In brief, the procedure iterates through each price series, starting from the second-to-last value and moving backward. If a price matches the final known price, we continue checking earlier values. Once a different price is found, all subsequent repeated values (i.e., those following the termination of active trading for the security) are replaced with "NA". This procedure ensures that natural price repetitions on non-trading days are preserved, while only artificial price extensions after delisting are removed.

The data quality is high, out of the 20,000 individual price series, we had to delete only about 200 due to various issues. These deletions occurred prior to running any analysis relevant to the objectives of this paper and the subset of removed stocks is small enough that it will not affect the overall results.

The primary goal of this paper is to examine questions related to the performance of the momentum factor under different conditions, as well as its predictability. However, we also incorporate additional data that requires further explanation, specifically the market capitalization and trading volume for the securities in our dataset. We collected this data on the first trading day of each month and then forward filled it for the remainder of the month. This approach was chosen to reduce the size of the large datasets we are working with. Since the momentum portfolios are rebalanced monthly, this method does not result in a loss of critical information and introduces only minimal inaccuracies when calculating market capitalization-weighted performance for each exchange. Additionally, we downloaded 3-month sovereign debt interest rates for the countries where our markets are located, which enabled us to calculate accurate Sharpe ratios and analyze the effect of interest rates on momentum performance.

4 Methodology

To evaluate whether different volatility forecasting models can meaningfully improve the performance of a dynamic momentum strategy, we first construct a baseline momentum portfolio using the approach of Jegadeesh and Titman (1993). The selected universe of stocks is sorted based on cumulative returns. The best and worst performing deciles of stocks are branded as winners and losers, winners are bought, and losers are sold. This also requires the definition of a lookback, formation and holding period. In alignment with "Momentum Crashes" by Daniel and Moskowitz (2016) the lookback period is set to 11 months, starting 12 months back from the formation date. The formation period of 1 month between the lookback and holding period allows the constructed portfolios to avoid the effects of short-term reversals as observed by Jegadeesh (1990) and Lehmann (1990). The holding period is 12 months following the formation date. In the appendix, as a robustness check, we also present results for a six-month version. Construction of a new momentum portfolio is repeated once every month, yielding a basket of 12 active portfolios at any time, except for the first and last 12 months of the studied period. Performance of the overall portfolio is calculated monthly as the aggregate return of all current portfolios.

Once the portfolio is constructed, we evaluate two approaches to momentum investing: a static buy-and-hold strategy and a dynamic weighting strategy. The simple buy and hold strategy continually holds 100% of the capital in the momentum portfolio, while the dynamic weighting strategy adjusts the allocation between a risk-free asset and the momentum portfolio. The rebalancing of the portfolio allocation is done at monthly intervals.

The weighting is decided using the methodology of Daniel and Moskowitz (2016), which incorporates forecasts of expected returns and volatility. They show that the function for the optimal weight takes the following form:

$$w_{t-1}^{*} = \left(\frac{1}{2\lambda}\right)\frac{\mu_{t-1}}{\sigma_{t-1}^{2}} \tag{1}$$

where μ_{t-1} is the conditional expected return on the momentum portfolio, σ_{t-1}^{2} is the conditional variance of the momentum portfolio and λ - lambda is a scalar used to control the exposure of the portfolio. The scalar can be adjusted to reach a certain level of portfolio volatility and enables comparison to static momentum portfolios.

To perform an analysis that is grounded in a more practically applicable methodology, we constrain the weights of our dynamic momentum portfolio to lie within the range [0, 1]. In contrast, Daniel and Moskowitz's implementation of the dynamic weighting framework allows optimal weights to exceed 1 or fall below 0, thereby enabling leveraged long or short positions in the momentum portfolio. While such flexibility may enhance performance, it reduces the strategy's practical applicability, particularly for investors subject to leverage constraints, margin requirements, or regulatory limitations. By imposing bounds on portfolio weights, we increase the practical relevance of our analysis and ensure it better reflects real-world investment considerations.

We estimate the conditional expected return μ_{t-1} , using the same regression method as Daniel and Moskowitz (2016), given by the following specification:

$$\underline{R}_{WML,t} = \gamma_0 + \gamma_{B,t-1} * I_{B,t-1} + \gamma_{\sigma_m^2} * \hat{\sigma}^2{}_{m,t-1} + \gamma_{int} * I_{B,t-1} * \hat{\sigma}^2{}_{m,t-1} + \underline{\epsilon}_t \qquad (2)$$

The momentum portfolio returns are modeled as a function of an interaction term, γ_{int}, between the market variance over the preceding 126 trading days, $\hat{\sigma}^2_{m,t-1}$, and a bear-market indicator, $I_{B,t-1}$. The indicator is a measure of cumulative returns of the market in the previous two years and equal to 1 if they are negative.

Our approach to determining the lambda parameter differs from that of Daniel and Moskowitz. While they select lambda to match the annualized portfolio volatility to a target level using full-sample information, an approach that involves ex-post data, we implement a strictly ex-ante method. Specifically, we re-estimate lambda at each monthly rebalancing point using a bounded optimization procedure. This procedure uses a rolling five-year window of historical data to identify the lambda value that would have aligned the portfolio's volatility with that of the market over that period. The estimated lambda is then applied to the subsequent month's portfolio allocation, and the process is repeated the next month. We use market volatility as the benchmark as this is an acceptable level of volatility for most investors, though our approach allows one to target any level.

Central to the analysis, this paper extends the volatility forecasting component by exploring multiple model specifications rather than relying on a single approach. We use six different models to generate portfolio volatility forecasts. These forecasts are then individually applied within the dynamic weighting function, resulting in a set of six distinct momentum strategies,

each corresponding to a different volatility model. The performance of these dynamic portfolios is evaluated and compared using return metrics and Sharpe ratios, providing insight into the effectiveness of each volatility forecasting method.

The range of model specifications applied was selected because it provides a broad evaluation in terms of model complexity and performance. Previous literature incorporating volatility forecasting for momentum strategies varies in relation to simpleness of application. First is the question of constant or dynamic volatility portfolios. A constant volatility momentum strategy adjusts the portfolio allocating to match a certain level of volatility. Whereas a dynamic volatility portfolio allows the level of volatility to be time-varying. Barroso and Santa-Clara (2015) implement a constant volatility strategy using a simple realized variance method to forecast volatility. This paper offers an easily implementable strategy which does not require sophisticated volatility-modelling. Daniel and Moskowitz (2016) fall on the other end of the spectrum with their dynamic volatility strategy utilizing a GJR-GARCH model for volatility forecasting, as well as a regression model to forecast expected returns. The set of models evaluated here will offer insight into the tradeoff relationship between forecast accuracy and complexity and performance gain expressed as the Sharpe ratio.

4.1 Models

The simplest model in terms of implementation is a realized variance model. This is included to provide a benchmark as a computationally simple alternative that requires no assumptions about the time-series characteristics. It is defined as

$$\hat{\sigma}^2_{WML,t} = 21 \sum_{j=0}^{125} \frac{r^2_{WML,d_{t-1-j}}}{126} \tag{3}$$

where an equal-weighted average of the daily squared returns over the last 6 months are multiplied by 21 to get a forecast for next month's variance level.

The next model is a baseline GARCH (1,1) model defined as

$$\sigma^2{}_t = \omega + \alpha \cdot \varepsilon^2{}_t{}^{-1} + \beta \cdot \sigma^2{}_t{}^{-1} \tag{4}$$

here, σ^2_t represents the conditional variance of momentum returns at time t, ε_{t-1} is the lagged residual (return shock), ω is a constant term for the long-run average variance, α is the weight on the lagged squared return shock (innovation term) and β is the weight on the lagged variance (persistence term).

We test several extensions of the GARCH (1,1) model, starting with the exponential GARCH (EGARCH) model. This version allows for asymmetric effects from positive versus negative shocks through a parameter θ. As discussed previously and demonstrated by Daniel and Moskowitz (2016), momentum portfolios exhibit volatility that reacts asymmetrically to return shocks, especially during sharp reversals. EGARCH has the potential to capture this feature better than standard GARCH. It is defined as

$$\ln(\sigma^2{}_t) = \omega + \beta \cdot \ln(\sigma^2{}_t{}^{-1}) + \gamma \cdot \left(\frac{|\varepsilon_t{}^{-1}|}{\sigma}{}_t{}^{-1} \right) + \theta \cdot \left(\frac{\varepsilon_t{}^{-1}}{\sigma}{}_t{}^{-1} \right) \qquad (5)$$

and differentiates from the baseline GARCH model by, $\ln(\sigma^2{}_t)$ which represents the logarithm of the conditional variance of momentum returns at time t, β is now the weight on the logarithm of past variance, γ captures the symmetric response of volatility to new information while θ captures the asymmetric response of volatility to shocks and $(\varepsilon_{t-1}/\sigma_{t-1})$ represents the standardized return shock.

Another specification of the GARCH model is the GJR-GARCH. Similarly to the EGARCH it allows for asymmetric response to shocks which is an observed characteristic of the momentum portfolio. By including an indicator function it captures the differential impact of bad news on volatility. It is defined as

$$\sigma^2{}_t = \omega + \alpha \cdot \varepsilon^2{}_t{}^{-1} + \gamma \cdot \varepsilon^2{}_t{}^{-1} \cdot I\{\varepsilon_t{}^{-1} < 0\} + \beta \cdot \sigma^2{}_t{}^{-1} \qquad (6)$$

here, γ captures the additional effect of negative return shocks on volatility and $I\{\varepsilon_{t-1} < 0\}$ is an indicator function that equals 1 when ε_{t-1} is negative.

To test whether macroeconomic variables improve volatility forecasts, the models above can be augmented with exogenous regressors. To test this, a GARCH-X model will also be fitted. It is defined as

$$\sigma^2{}_t = \omega + \alpha \cdot \varepsilon^2{}_t{}^{-1} + \beta \cdot \sigma^2{}_t{}^{-1} + \delta \cdot X_t{}^{-1} \qquad (7)$$

where δ measures the influence of external variables on the conditional variance and X_{t-1} is a vector of lagged exogenous variables, in this case the three-month treasury or gilt interest rate.

Finally at the end of the spectrum in terms of model complexity an ARIMA-GARCH model is tested. Here, the best performing GARCH model in terms of forecasting precision will be combined with an ARIMA model. The ARIMA model is included to capture the conditional mean process of the momentum portfolio returns, accounting for autocorrelation, trends, and

other structural features in the return series. The residuals from this mean equation are then used as inputs for the GARCH model, which models the conditional variance of the returns. The ARIMA-GARCH model's relevance to the momentum portfolio volatility stems from the existence of potential serial dependence in returns and distinct volatility patterns following periods of market stress or reversal. The ARIMA model is defined as

$$\mu_t = c + \sum_{i=1}^{p} \phi_i r_{t-i} + \sum_{j=1}^{q} \theta_j e_{t-j} \tag{8}$$

where μ_t is the expected return at time t, c is a constant representing the drift in the return process. The autoregressive terms, $\sum_{i=1}^{p} \phi_i r_{t-i}$, are the lagged values of the return series where p is the number of lags and ϕ_i are the coefficients on past returns indicating the level of influence on expected returns. The moving average terms, $\sum_{j=1}^{q} \theta_j e_{t-j}$, capture the lagged forecast errors where q is the number of lags and θ_j are the coefficients on past shocks measuring the impact on expected returns.

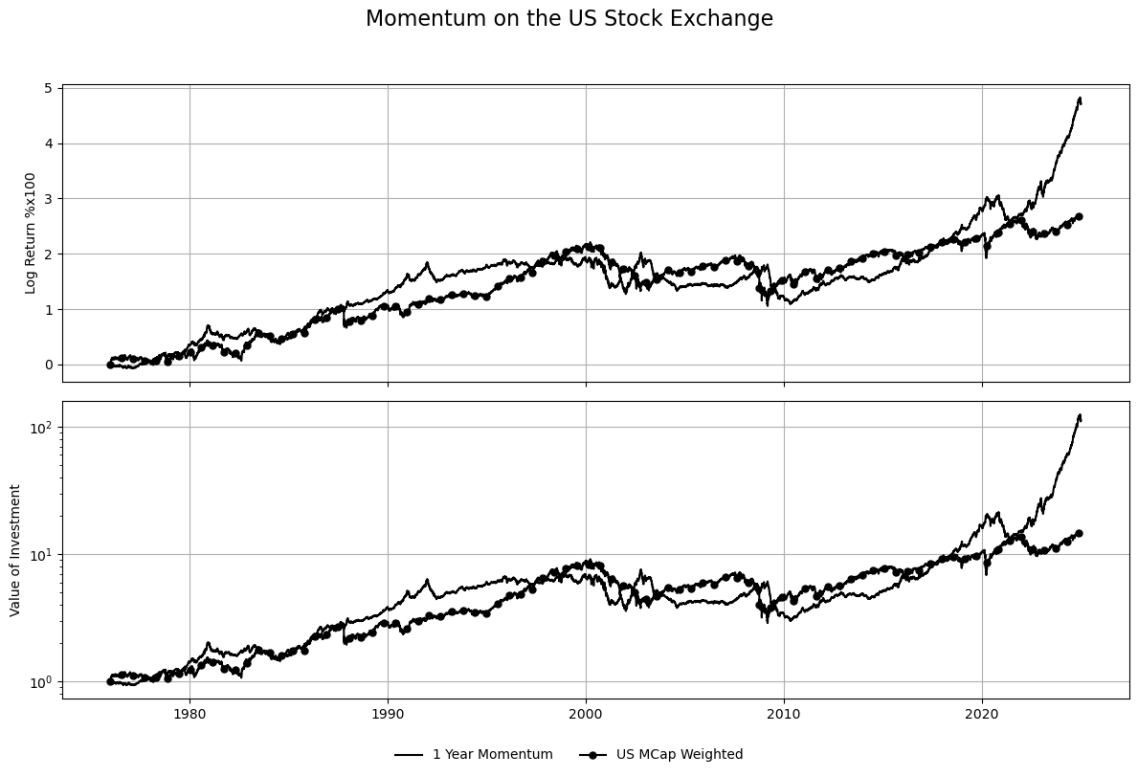
For each of the GARCH models, parameters are estimated using maximum likelihood estimation. At each monthly rebalancing point, the estimation is performed using data from the preceding five years, ensuring that only information available up to that point is used.

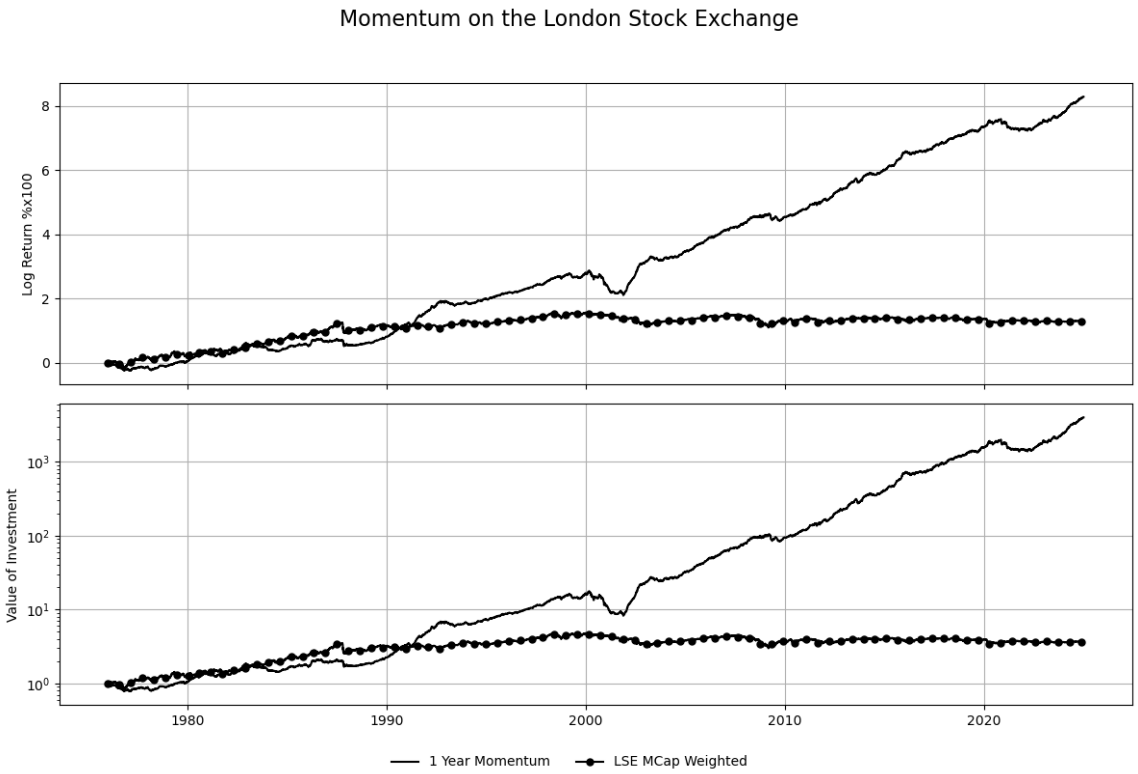
5 Analysis of the Momentum Performance

The performance of the momentum factor has been extensively researched since its initial discovery by Jegadeesh and Titman (1993). Rather than offering a purely normative description in this section, we aim to highlight stylized facts that support the methods used to address the main research question of this paper. Moreover, since our sources date from 2004 to 2015, they do not account for the notably eventful second half of the 2010s, particularly the Covid-19 pandemic followed by the sudden rise of global inflation. These developments led to unprecedented market movements, which have been scarcely examined from the perspective of momentum performance.

Momentum consistently generates significant and robust returns across a wide range of assets and configurations. In this section, we first present the performance characteristics of the momentum strategy over the past decade, followed by an updated analysis of Cooper, Gutierrez, and Hameed (2004). This approach aims to further support the findings referenced above.

The first set of charts illustrates the momentum factor's performance on each exchange since 1975. We focus on the one-year momentum, which refers to two key periods: first, the time we look back from the formation date to observe the performance of the stocks used for sorting, and second, the duration for which we hold the portfolio after the formation date. For reference, we also show the market capitalization-weighted performance of all stocks on each exchange. The momentum strategy returns are calculated by including every available stock on the exchange, without any further filtering unless explicitly stated.





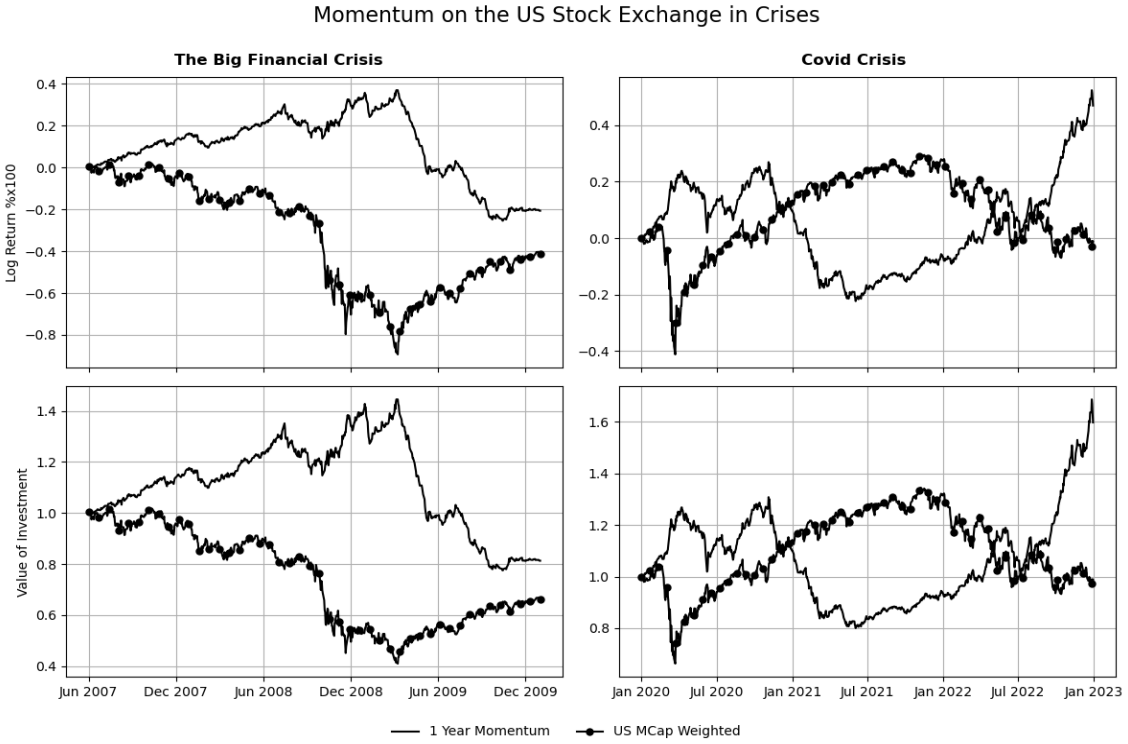
Figures 2 – 3. Overview of Momentum Performance on the Combined U.S. Market (Nasdaq and NYSE) and the London Stock Exchange, from 1976 to 2024. We show the one-year momentum, with both the formation and holding periods lasting one year. For comparison, we include the market capitalization-weighted performance of all stocks within the index. Additionally, both charts illustrate the value of an initial investment of 1 unit.

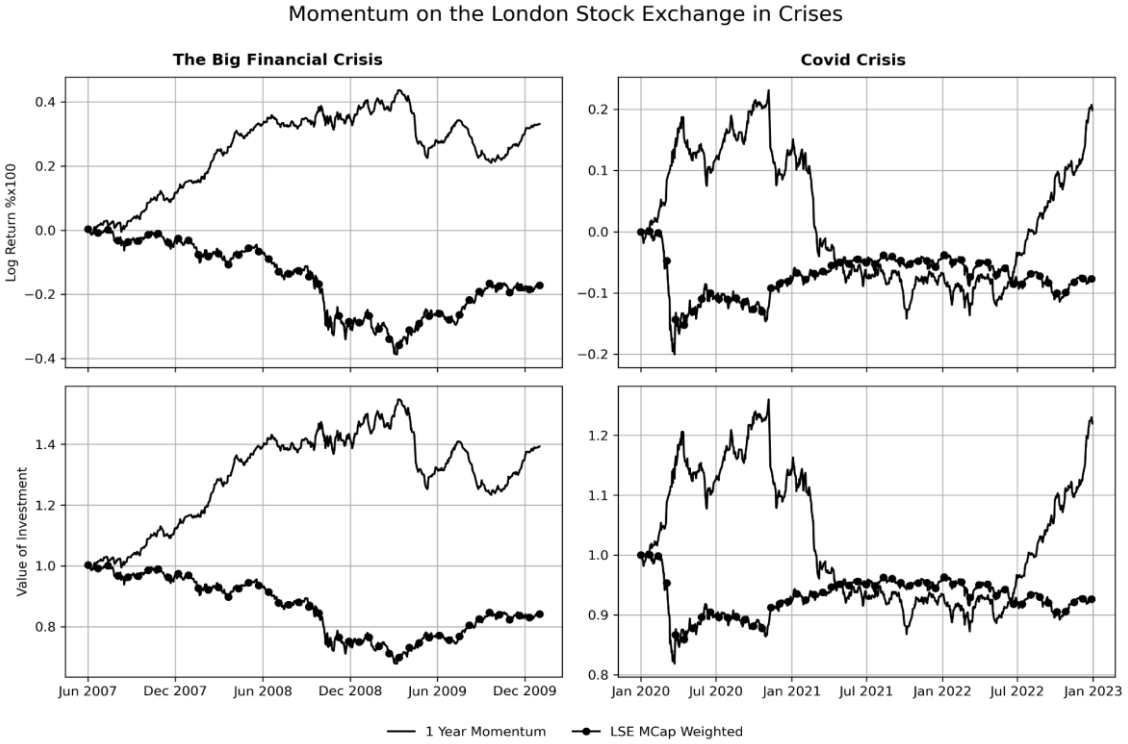
Table 1. Overview of Momentum Performance in the Analyzed Markets. Using the respective country specific annualized return on 3-Month sovereign debt securities as the risk-free asset, from 1976 to 2024

Market	Type	Annual Excess Return	Annual Standard Deviation	Sharpe Ratio
US	Market	2.3%	16.8%	0.14
LSE	Market	-1.6%	8.9%	-0.18
US	Momentum	6.5%	11.0%	0.59
LSE	Momentum	12.9%	8.2%	1.58

The momentum factor generally outperforms its respective market environment, even over the long run. However, momentum on the NYSE and Nasdaq only slightly outperforms its respective market index. In the last decade, the long-term momentum strategy continues to generate positive returns relative to the broader market.

Daniel and Moskowitz (2016) pays particular attention to crisis periods, noting that momentum strategies tend to perform poorly following market crashes. They propose several potential explanations, including the option-like payoff of recent losers and the time-varying exposure of momentum strategies to volatility risk. This phenomenon persists after their observation period, which ended in 2013.

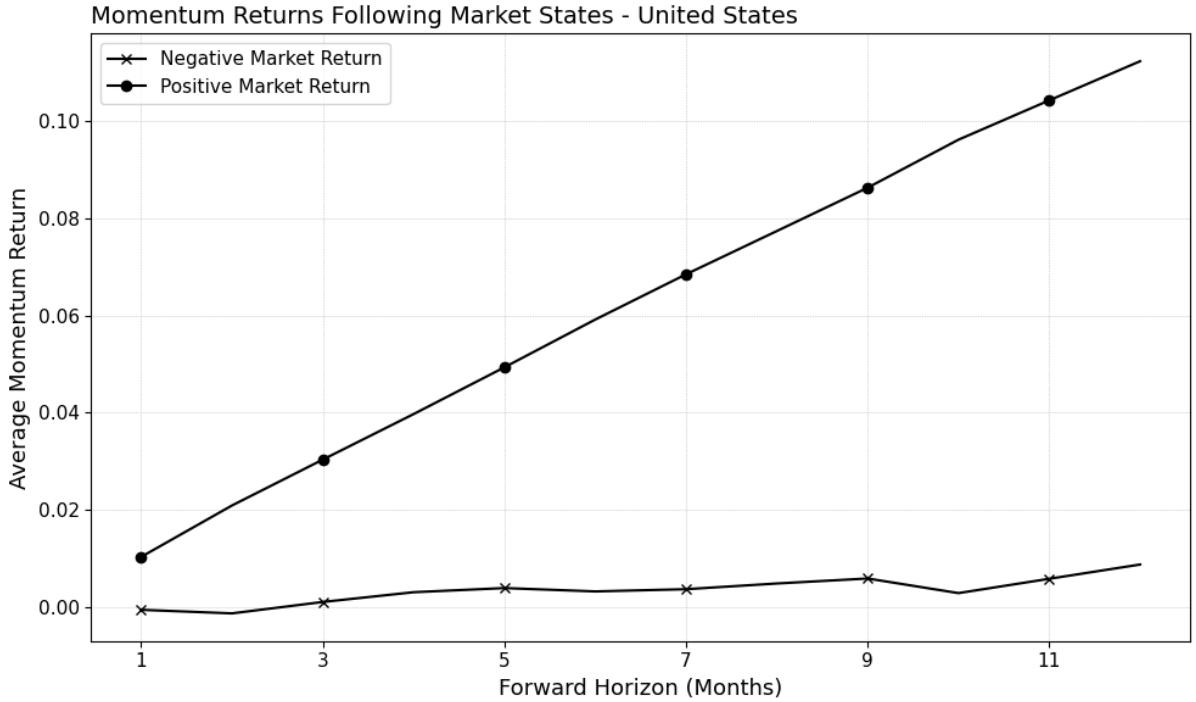


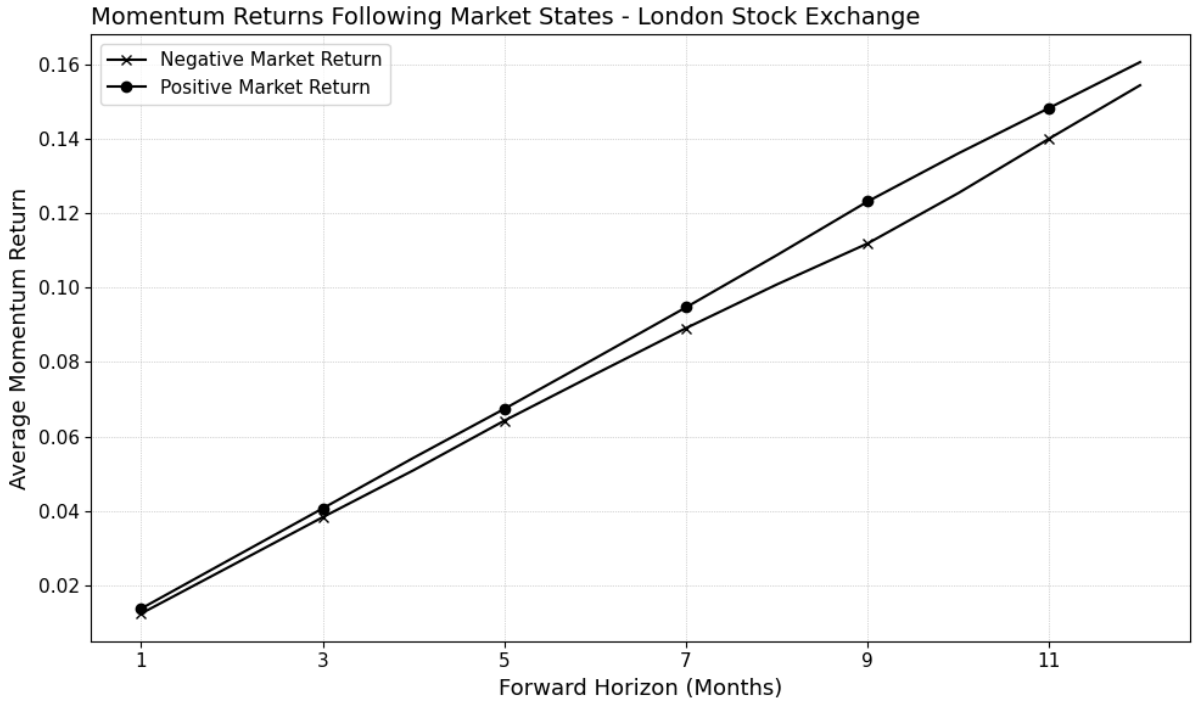


Figures 4 – 5. Overview of Momentum Performance During Crisis Periods. We compare the performance of the one-year momentum portfolio on both markets with the market capitalization-weighted performance of all stocks within the index. The chart specifically focuses on the last two major global financial crises and illustrates the challenge momentum strategies face when the broader market is rebounding.

The market rebound-momentum loss phenomenon is prevalent during both the 2008 financial crisis and the 2020 Covid-19 crisis. These results are robust across the markets we examined. Our data exhibits the same characteristics as that used by Daniel and Moskowitz (2016). This pattern is not a one-time occurrence; rather, it is a consistent feature of momentum strategies. Therefore, there is a clear motivation to address the time-varying exposure to volatility risk in momentum strategies, as these events represent the primary and almost exclusive loss leaders of the trading rule.

In the final step, we replicate the analysis conducted by Cooper, Gutierrez, and Hameed (2004) using our data and an extended time frame. The authors find that the performance of the momentum factor is significantly better following positive market performance. We employ a similar algorithm: we assess market performance over the past three years and assign a "bucket" based on the result. An "UP" bucket is designated if the market has experienced positive performance, and a "DOWN" bucket if the market has not. We then analyze the momentum strategy's results following these time frames, recording observations monthly. By looping through the entire dataset, we calculate conditional averages based on the buckets to which the observations belong.





Figures 6 – 7. Momentum Strategy Performance Conditional on Preceding Market States. We examine how the conditional average returns of a one-year momentum strategy evolve over a 12-month horizon, depending on whether the market closed positively in the previous three years. Conditional underperformance is notably more pronounced in the U.S. markets. In contrast, the momentum strategy on the London Stock Exchange performs only slightly worse when the market closed with a loss.

The charts clearly confirm the original observations of Cooper, Gutierrez, and Hameed (2004). These findings remain valid even after 20 years, which enhances their implied statistical significance. Momentum performs strongly following periods of upward market trends, particularly in the U.S. In contrast, the UK exhibits only a slight outperformance. Nevertheless, these findings continue to hold true over an extended period and across a broader geographic scope.

The descriptive characteristics of the momentum strategy remain empirically valid decades after their initial observation. More specifically, this section highlights the issue of time-varying exposure to volatility, which continues to be a concern. Consequently, developing strategies to mitigate this volatility remains highly relevant.

6 Results

This section presents the empirical findings of the study. We begin by reporting the main results for the static momentum strategy and six dynamic strategies, using metrics for both forecast accuracy and portfolio performance. From an economic perspective, portfolio performance is assessed in terms of excess returns, realized volatility, and Sharpe ratios. Forecasting accuracy is evaluated using root mean squared error (RMSE) and mean absolute error (MAE), consistent with previous work such as that of Hansen and Lunde (2005).

We then provide visual representations to further illustrate the comparative forecast efficiency of the models. This is followed by a detailed breakdown of the dynamic momentum strategy implementation for two different model approaches, the realized variance and the simple GARCH (1,1). These two are selected as benchmarks given that the broader implementation patterns do not differ substantially across models.

Finally, special attention is paid to the performance of the dynamic strategies during periods of market stress.

Table 2.
Volatility forecast accuracy and strategy performance.
These tables present forecast accuracy measures (RMSE and MAE) and portfolio performance metrics for the static momentum strategy and six dynamic strategies. Forecast accuracy metrics are computed in daily return units, while excess returns, volatility, and Sharpe ratios are annualized. Volatility predictions are scaled to match realized volatility. Panel A presents results for the US market (Nasdaq and NYSE) and Panel B presents results for the London Stock Exchange. Note: Return measures are based on price data only and do not include dividends. This may affect comparability with studies using total return data but does not affect the results of this analysis.

Panel A.

US market.

Model	RMSE	MAE	Annual Excess Return	Annual Volatility	Sharpe-Ratio
Static Momentum	n/a	n/a	4.74%	10.96%	0.4320
Realized Variance	0.000953	0.000597	5.00%	9.58%	0.5219
GARCH	0.000364	0.000208	4.88%	9.62%	0.5067
EGARCH	0.002130	0.000984	5.01%	9.58%	0.5237
GJR-GARCH	0.000292	0.000161	4.90%	9.62%	0.5095
X-GARCH	0.001351	0.000809	4.69%	9.67%	0.4853
ARIMA-GARCH	0.003569	0.002817	4.79%	9.54%	0.5018

Panel B

London Stock Exchange.

Model	RMSE	MAE	Annual Excess Return	Annual Volatility	Sharpe Ratio
Static Momentum	n/a	n/a	13.17%	8.24%	1.5988
Realized Variance	0.000867	0.000633	11.64%	7.11%	1.6380
GARCH	0.000139	0.000093	11.96%	7.21%	1.6591
EGARCH	0.001704	0.000728	11.73%	7.18%	1.6332
GJR-GARCH	0.000142	0.000091	11.92%	7.20%	1.6545
X-GARCH	0.001303	0.000805	11.40%	7.38%	1.5450
ARIMA-GARCH	0.003174	0.002526	11.38%	7.39%	1.5387

Table 2 presents the main results of the seven different momentum strategies. Panel A shows results for the US market and panel B for the LSE. We further evaluated our results by applying a 6-month sorting and holding period strategy, see table 5 in the appendix.

As seen in table 2 the simple GARCH and GJR-GARCH generate the most accurate forecasts across both markets, as measured by RMSE and MAE. Additionally, the realized variance model performs consequently well across the two markets and is the third most efficient predictor, beating out the three less efficient GARCH-based models. The ARIMA-GARCH model is the least efficient predictor based on these results.

In terms of portfolio performance, Table 2 also shows that the static momentum strategy and the six dynamic strategies based on different volatility models deliver varied results. Overall, portfolios in the London market delivered higher excess returns compared to those in the US, with the US portfolios also exhibiting greater volatility. These results are simply a consequence of the stock universes utilized in this paper. Since the number of stocks is large and these companies vary widely in terms of size and other characteristics, no concrete conclusions should be made about the relative size of returns. The methodology of this paper is designed to facilitate the comparison of volatility handling and modeling, rather than to draw conclusions about overall momentum return performance across different markets.

When compared to the static momentum strategy, the dynamic momentum strategies in the US outperformed by generating both higher excess returns and lower volatility. In contrast, the dynamic strategies in the London market did not exceed the excess return of the static strategy. However, four models in the London market were able to reduce volatility enough to achieve a higher Sharpe ratio, nonetheless.

Panel A shows that the EGARCH-based dynamic approach performed best in terms of Sharpe ratio in the US, narrowly beating the realized variance approach. For the London market, the basic GARCH-based approach produced the highest risk-adjusted returns. Because these numbers differ only marginally, any deterministic conclusions about the ordering of performance would be statistically unsubstantiated.

Interestingly, the results in Table 2 indicate that higher forecast accuracy does not necessarily lead to better risk-adjusted performance. For example, in Panel A, although both EGARCH and the realized variance model have lower forecasting accuracy than GARCH and GJR-GARCH, they achieve a higher Sharpe ratio. This disconnect between volatility prediction accuracy and portfolio performance appears in several instances.

As expected, due to the lambda parameter in the weighting function being used to match market volatility, the volatility levels are roughly equal for all models in both markets. The dynamic approach decreases volatility in contrast to the static momentum strategy regardless of which volatility model that is used.

6.1 Volatility forecasts

Taking a more detailed look at volatility model forecasts, Figure 8 displays the 21-day rolling predicted volatility for the three most accurate models, along with the realized volatility. With a focus on two volatile periods, the 2008 financial crisis and the COVID crisis, these plots offer an insight into the characteristics of the different forecasts during periods of certain interest.

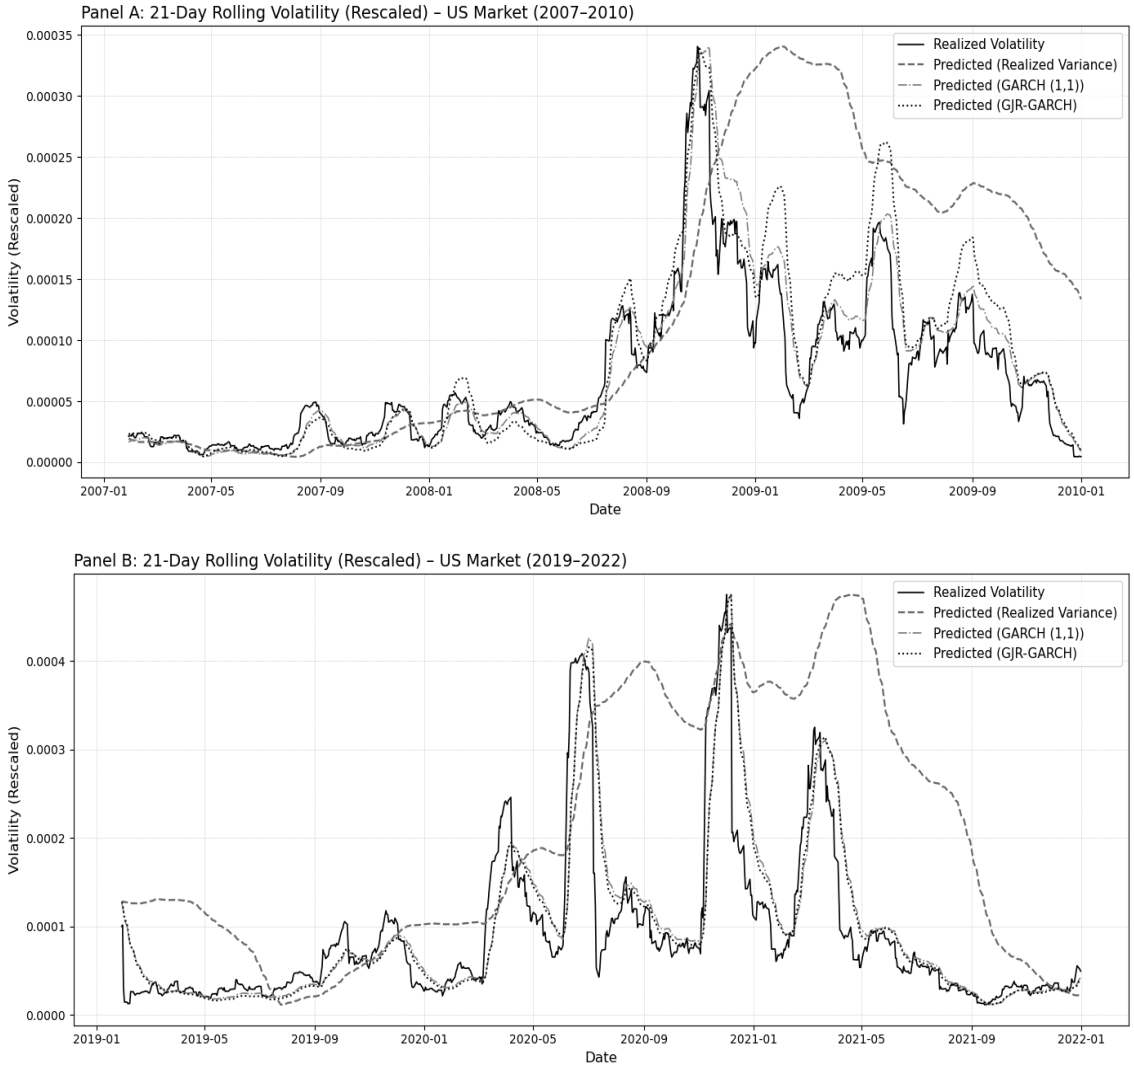


Figure 8. This figure displays realized and predicted 21-day rolling average volatilities for the realized variance, GARCH (1,1) and GJR-GARCH models, for the US market. The volatilities have been scaled for ease of comparison. Panel A presents the period surrounding the 2008 financial crisis. Panel B presents the period surrounding the COVID-crisis.

As seen in both panels, the realized volatility produces smoother forecasts relative to the GARCH models that fit the realized volatility more tightly and are faster to react to spikes and drops. In panel A we can observe that the GJR-GARCH predicts significantly higher spikes in volatility after the initial peak in 2008, compared to the basic GARCH. This is expected as the GJR-GARCH is constructed to react more aggressively to negative shocks, which the momentum portfolio experiences in the wake of the financial crisis.

6.2 Dynamic strategy performance

To examine the overall effectiveness of the different models within the dynamic strategy framework, figures 9 and 12 present plots for the cumulative returns, predicted volatility and momentum portfolio weight, for both the realized variance and GARCH (1,1) based dynamic strategies. The results are visually similar across all models, making individual presentations of all six model results redundant.

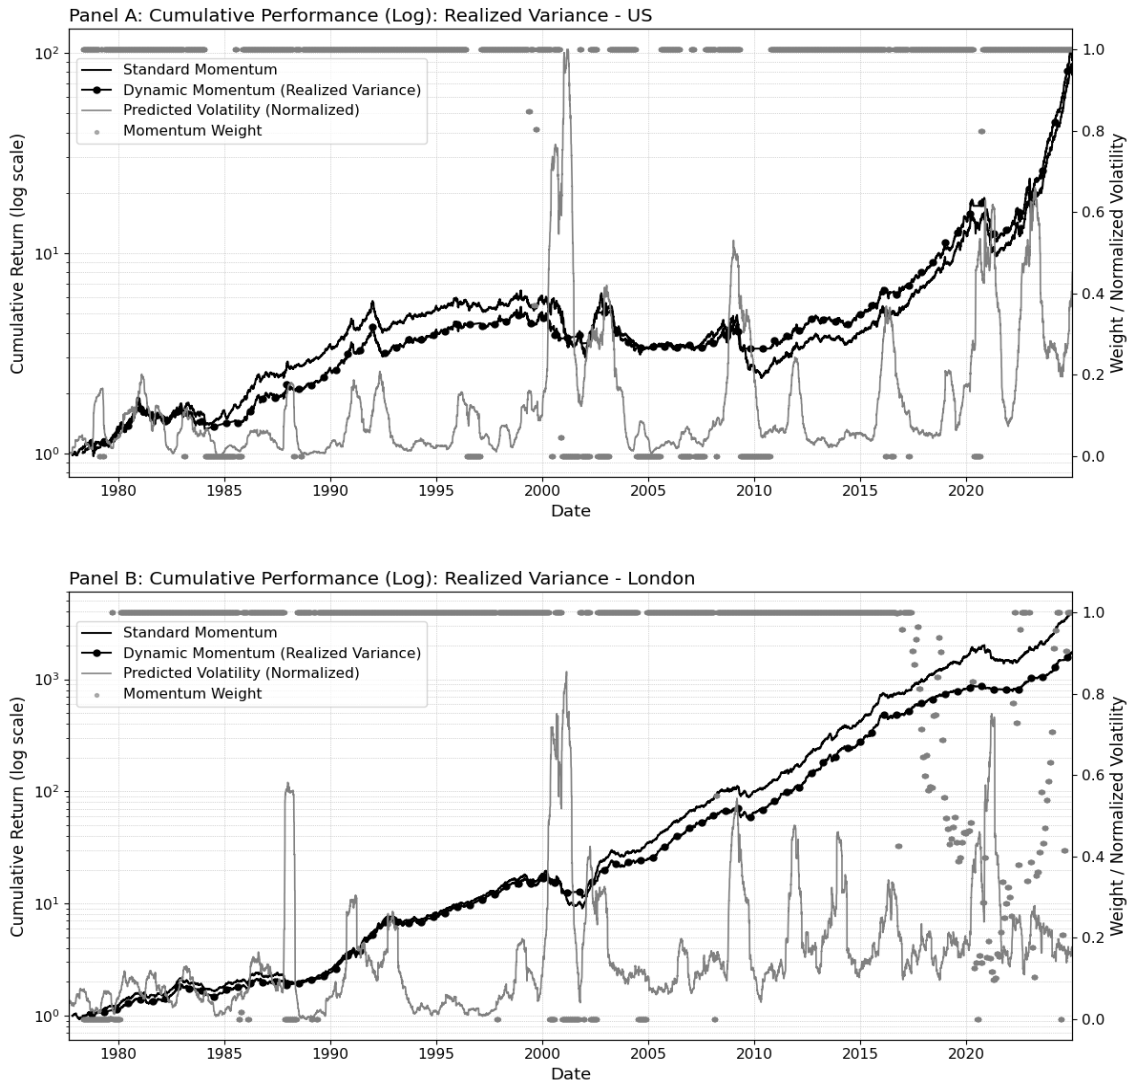


Figure 9. This figure displays the cumulative returns (on a logarithmic scale) of a standard momentum portfolio and a dynamically weighted momentum portfolio. The dynamic portfolio incorporates predicted volatility based on a realized variance model, with corresponding optimal portfolio weights also shown. Panel A presents results for the U.S. stock markets, while Panel B illustrates outcomes for the London Stock Exchange.

As shown by Figure 9, the realized variance approach yields returns that are broadly comparable to the static momentum strategy over the observed period. In the case of the US stock universe, the dynamic portfolio even achieves slightly higher cumulative returns. There, the static strategy outperformed between 1985 and 2000, a period dominated by steady gains. During the 2000's the dynamic strategy overtook in terms of cumulative returns. This is due to a shift to a risk-free allocation in the wake of the dotcom and 2008 financial crisis, avoiding subsequent drawdowns in momentum returns. For the London stock universe, the two strategies perform similarly. With the static strategy gradually producing better returns after the millennium shift and onward. Here, the turbulence of the 2000's is not as pronounced and rather, both strategies continue to yield positive returns.

Both Panel A and Panel B illustrate that the dynamic strategy performs as intended, it mitigates drawdowns during periods of heightened volatility. As shown in both panels, the portfolio allocation shifts to the risk-free asset in times of heightened predicted volatility, signaled by a flat return curve as the weights shift to zero. This effect is particularly pronounced during the 2008 financial crisis in the US market, where the dynamic portfolio reduces exposure to the momentum strategy and shifts toward a risk-free allocation. This adjustment significantly dampens the negative returns observed for the static momentum portfolio during that period.

Throughout most of the sample, the realized variance approach results in a binary allocation, fully invested in either the momentum or the risk-free asset. However, during the latter part of the sample for the London Stock Exchange, the portfolio weights exhibit greater variation, fluctuating in the range between 0 and 1. The explanation for this regime shift for the weighting function stems from the lambda parameter which controls the elasticity of allocation weights. A closer analysis of this is presented below.

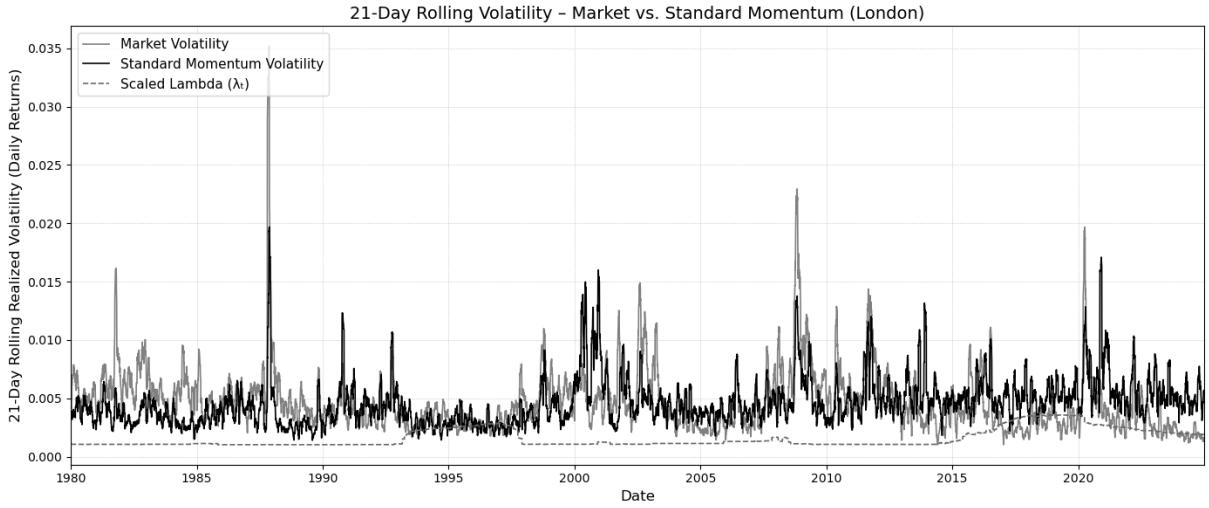


Figure 10. This figure plots the 21-day rolling volatility of the standard momentum portfolio and market portfolio for the London Stock Exchange along with lambda values scaled to fit the y-axis.

Since the weights are constrained between 0 and 1 and lambda is calibrated to match the momentum portfolio's volatility to that of the market, the volatility of the dynamic momentum portfolio is largely bound by the volatility of the standard momentum portfolio when the latter remains below the market portfolio's volatility. Figure 10 illustrates why there is a shift in the behavior of the dynamic weighting strategy from 2018 onward. During this period, momentum portfolio volatility begins to consistently exceed market volatility. Because the lambda parameter is calibrated to align the volatility of the dynamic strategy with that of the market, this change necessitates higher lambda values. A higher lambda reduces the responsiveness of

the portfolio weight to the expected return–to–volatility ratio, resulting in a more conservative allocation when predicted volatility is elevated.[1] This adjustment mechanism ensures that the dynamic strategy does not become excessively exposed when momentum returns are riskier than the benchmark.

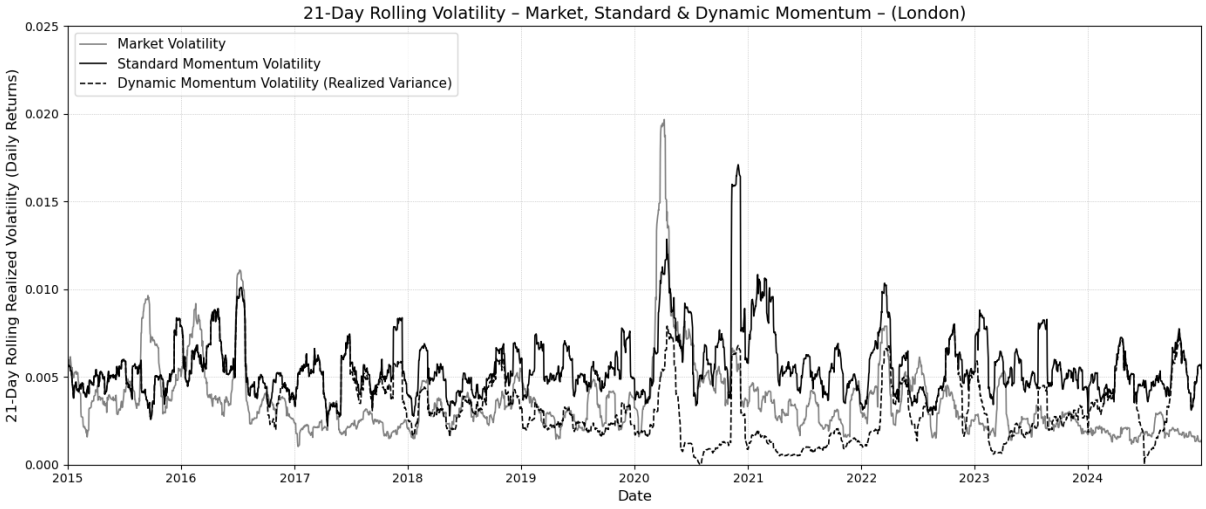


Figure 11. This plots the 21-day rolling volatility for the market portfolio, standard momentum portfolio and dynamic momentum portfolio, for the period 2015-2024 on the London Stock Exchange.

As shown in Figure 11, the lambda-scaling approach generally succeeds in aligning the volatility of the dynamic momentum portfolio with that of the market portfolio. This pattern is consistent across all six dynamic portfolio variations. Initially, when the standard momentum portfolio's volatility surpasses market volatility, the dynamic portfolio follows suit. However, starting around 2018, its volatility begins to diverge, fluctuating around the market volatility level until 2024, when it once again rises above it. The dynamic momentum volatility peaks significantly lower than both the market and standard momentum volatility during the COVID crisis, indicating that the dynamic weighting function is performing as intended.

[1] We present a detailed discussion about the behavior of Lambda in the Appendix.

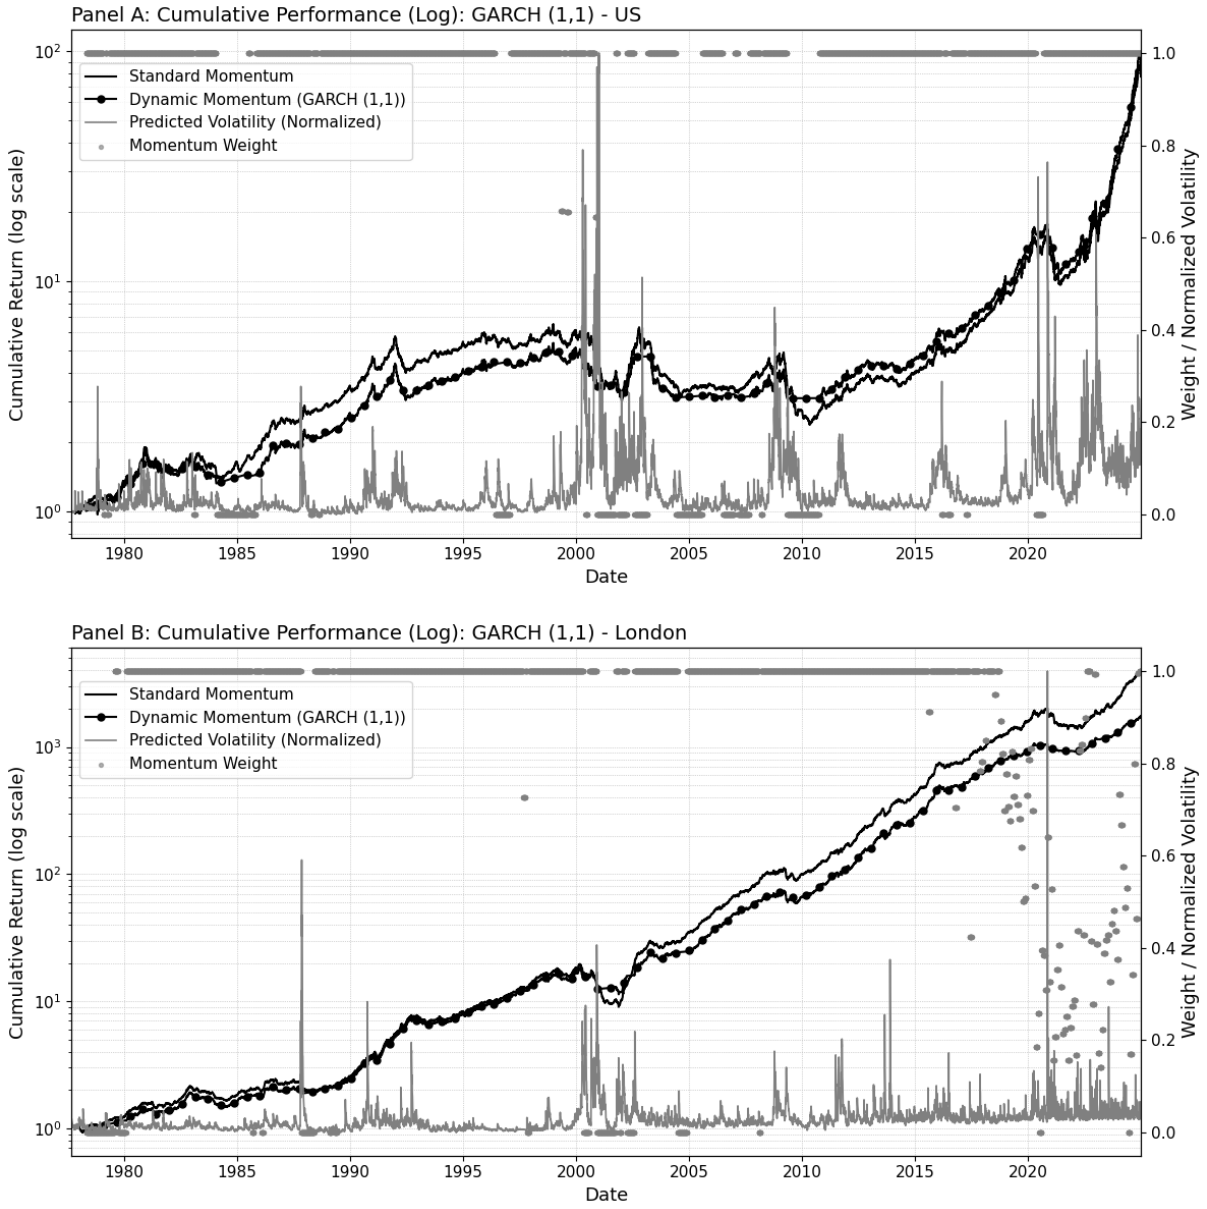
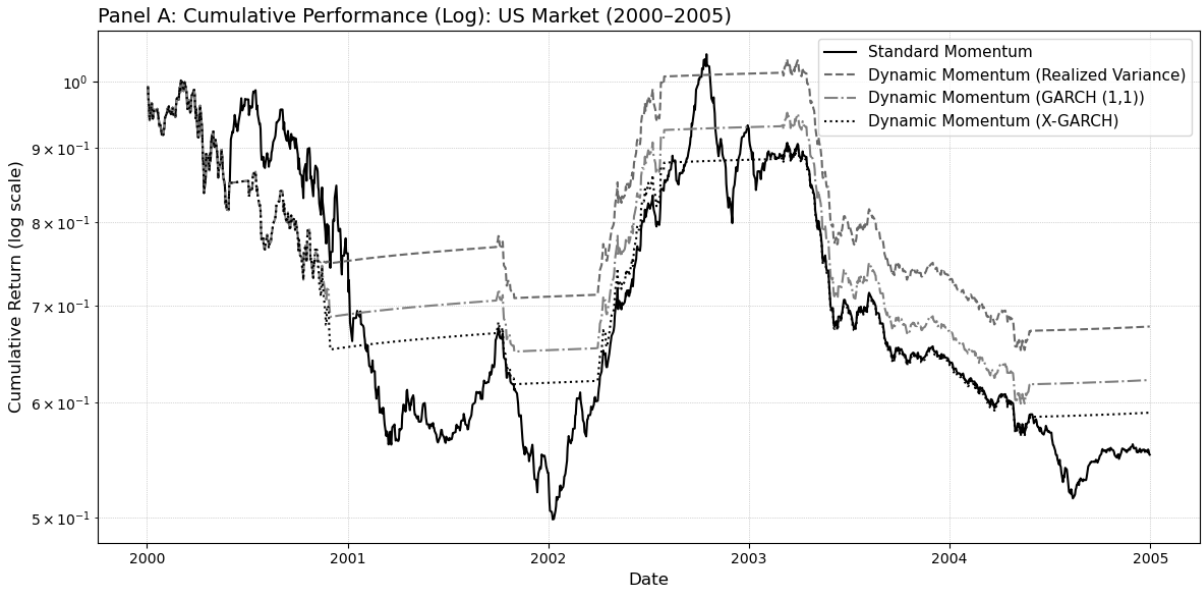


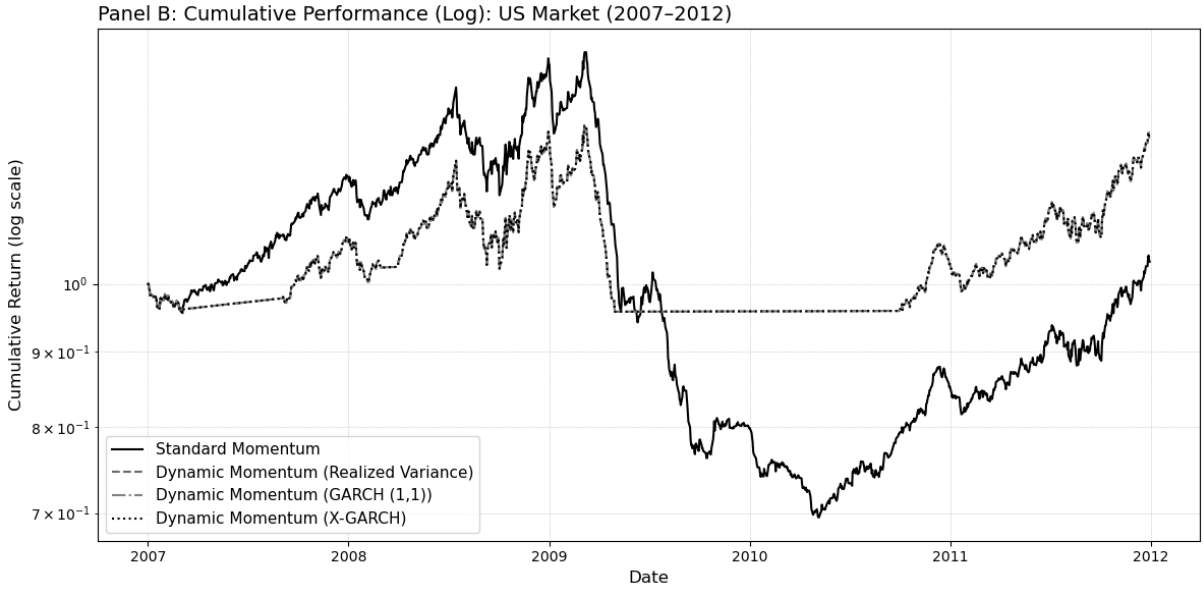
Figure 12. This figure displays the cumulative returns (on a logarithmic scale) of a standard momentum portfolio and a dynamically weighted momentum portfolio. The dynamic portfolio incorporates predicted volatility based on a GARCH (1,1) model, with corresponding optimal portfolio weights also shown. Panel A presents results for the U.S. stock markets, while Panel B illustrates outcomes for the London Stock Exchange.

In Figure 12, the same cumulative return plots are presented as in Figure 8, this time for the GARCH (1,1) dynamic strategy. Both panels visually confirm that its performance closely mirrors that of the realized variance-based strategy. As expected, predicted volatility rises during crisis periods, with more pronounced spikes observed in the U.S. market.

The cumulative return trajectories are largely consistent with earlier results. However, one notable difference is that the GARCH (1,1) strategy marginally underperforms the standard momentum strategy during peak periods around 2003 and 2009 in the U.S. This leads to a

smaller divergence in cumulative returns between the standard and dynamic strategies throughout the 2010s, compared to the realized variance-based approach. Finally, as observed previously, the weighting dynamics for the London market shift from approximately 2018 onward.





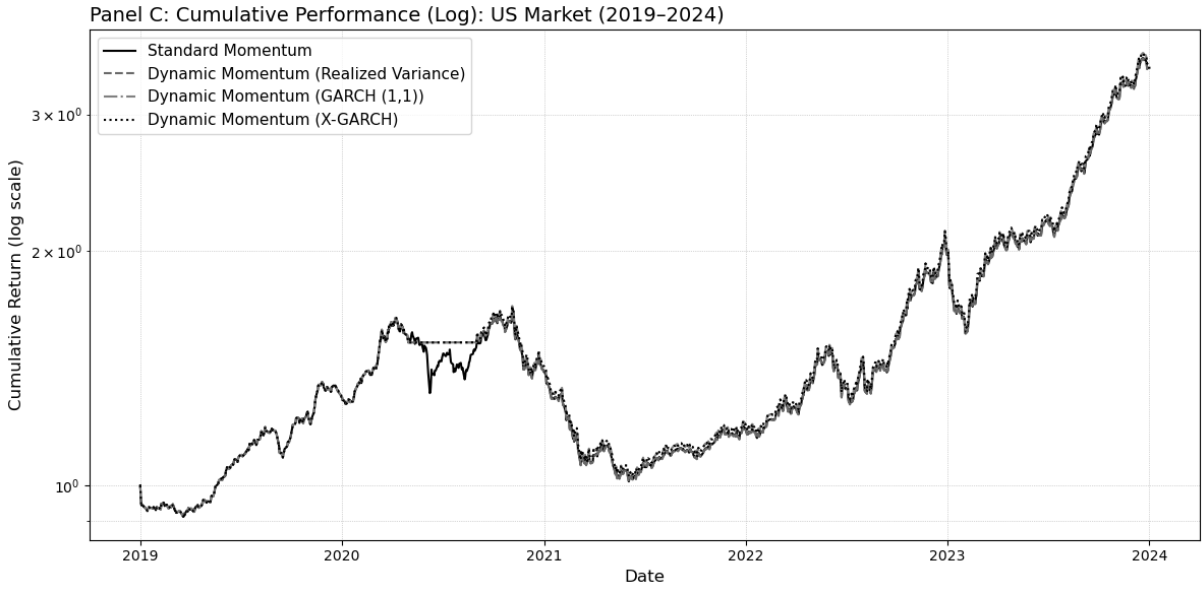


Figure 13. Panel A, B and C plot the cumulative performance (logarithmic scale) of the standard momentum, and three best performing dynamic strategies during periods of market turmoil in the US stock markets.

In Figure 13, the cumulative performance of three dynamic models is presented over three periods of market stress. These models were selected to highlight differences in performance across volatility models. The plots allow for analysis of any differences in the models' ability to manage turbulent periods and mitigate losses following market crashes, an issue central to the findings of Daniel and Moskowitz (2016).

Panels B and C show near-identical performance across the three dynamic strategies. In Panel B, all dynamic models initially underperform the standard momentum strategy but subsequently reverse this by shifting to a risk-free allocation during the post-financial crisis momentum crash. This mitigates what Daniel and Moskowitz (2016) point out as a key weakness of the momentum strategy. The tendency to suffer sharp losses following crashes, particularly when shorted "loser" stocks experience sharp reversals.

In contrast, Panel A reveals greater variation among the models. Here, the realized variance-based strategy delivers the highest cumulative returns. These differences are primarily driven by the timing of the transition to a risk-free allocation, which is governed by each model's volatility forecasts. The realized variance strategy is quicker to react to negative returns and increased volatility. Several periods of a complete allocation to the risk-free asset can be observed in panel A, which results in all three dynamic strategies yielding higher returns than the standard momentum strategy.

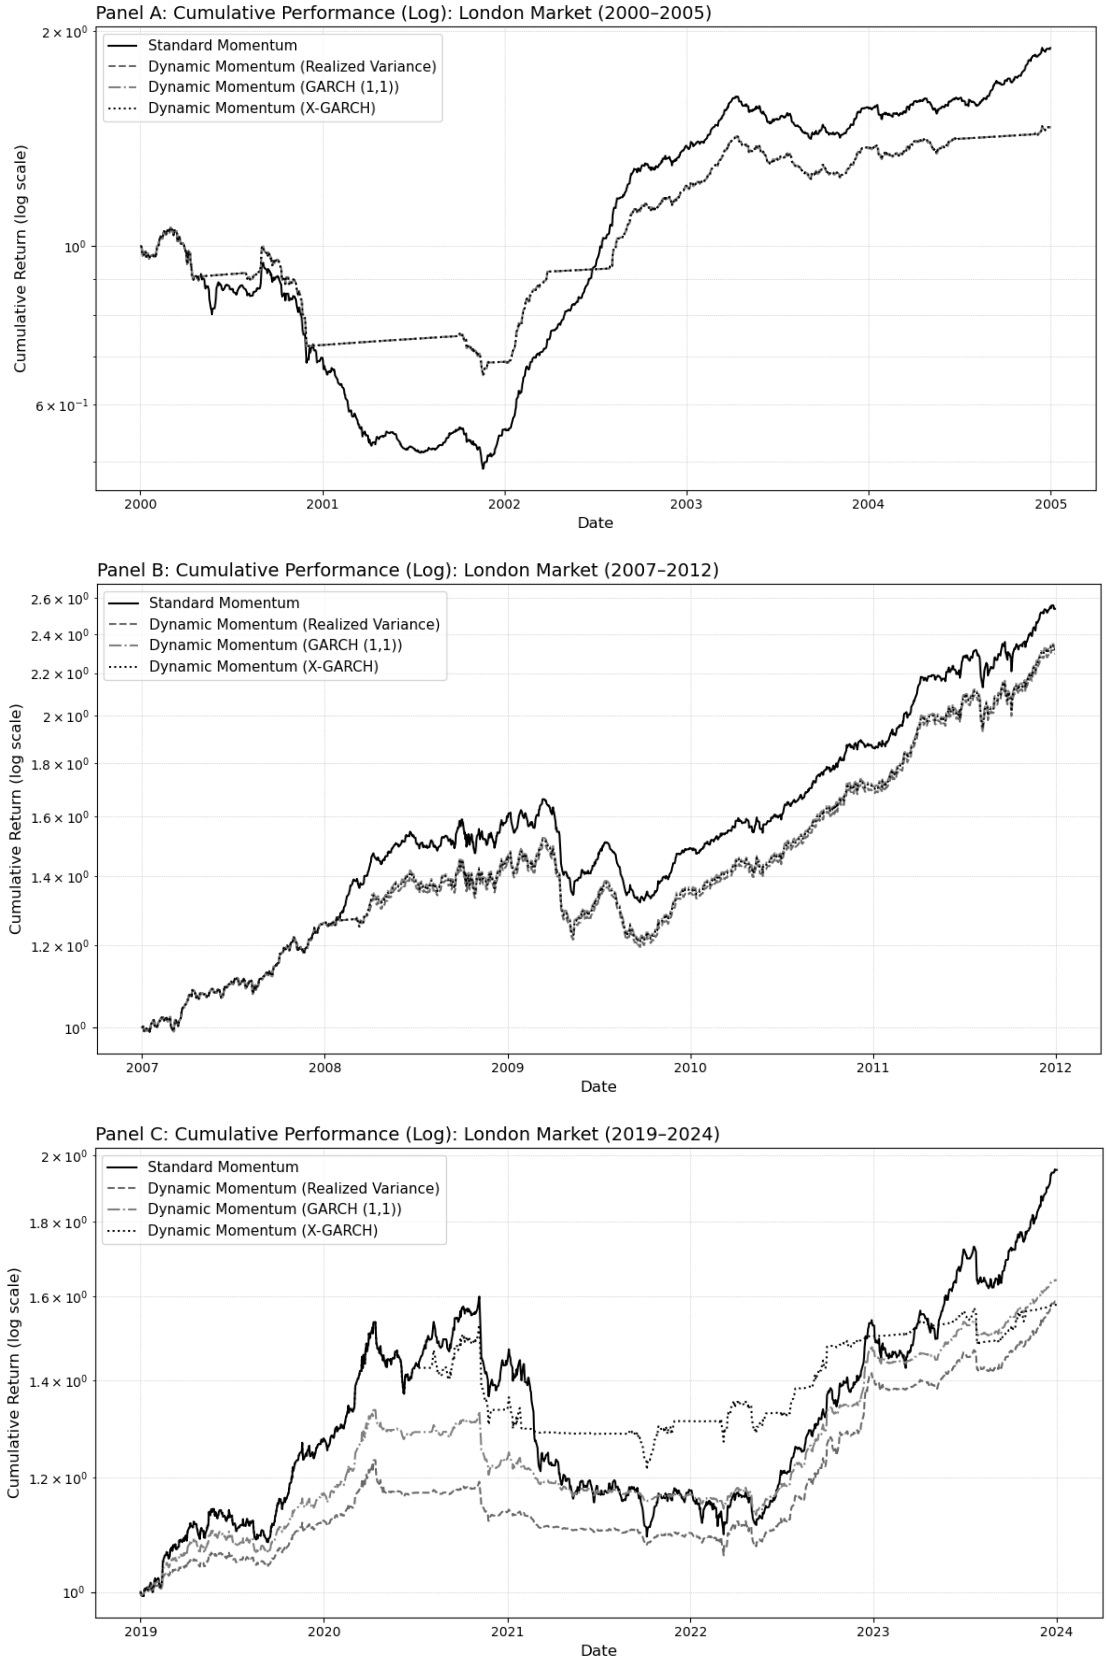


Figure 14. Panel A, B and C plot the cumulative performance (logarithmic scale) of the standard momentum, and three best performing dynamic strategies during periods of market turmoil on the London Stock Exchange.

Figure 14 presents the corresponding plots for the London Stock Exchange. Here, the three dynamic strategies exhibit nearly identical performance during both the dotcom bubble and the 2008 financial crisis. In each of the three periods, the dynamic strategies consistently underperform the standard momentum strategy in terms of cumulative returns.

Notably, the observed declines in returns do not coincide with the actual timing of the market crashes in Panels B and C. This suggests that, in the case of the London Stock Exchange, the momentum factor was not immediately affected by the 2008 financial crisis or the COVID-19 crisis. Rather, the delayed reaction aligns with the previously discussed phenomenon, where momentum strategies experience sharp reversals outside of the most acute phases of market turmoil.

During the COVID crisis, as shown in Panel C, the models once again diverge. The X-GARCH responds most effectively to minimize losses during late 2020 and early 2021. However, this same model proves less effective in capturing the strong rebound during the latter half of 2022, leading to a relative erosion of its earlier outperformance.

7 Discussion

The GARCH and GJR-GARCH models performed best in modeling the volatility of momentum portfolios. However, the realized variance model, a computationally simpler alternative, had comparable forecasting accuracy. These results suggest that, even when considering volatility forecasting in isolation, a simple volatility model may be sufficient for a dynamic momentum strategy.

Furthermore, the six dynamic strategies yielded similar risk-adjusted returns. In the U.S. stock universe, the realized variance and EGARCH models achieved the highest Sharpe ratios, despite a lower volatility forecasting accuracy than the basic GARCH and GJR-GARCH models. In contrast, for the LSE, the latter two models produced the highest Sharpe ratios, although the differences were marginal. Looking deeper into why higher forecasting accuracy does not necessarily translate into higher returns requires consideration of model characteristics. As presented in the results section, the realized variance model exhibits much higher persistence in the impact of past volatility shocks on forecasts of future volatility compared to the GARCH models. Consequently, this leads to an increased and more sustained allocation to the risk-free asset. This characteristic can be beneficial to a risk-averse investor unwilling to bear the additional risk of the momentum factor during high-volatility periods. In contrast, models with more reactive volatility estimates may better suit investors willing to accept higher portfolio turnover and volatility in exchange for potential performance gains. Thus, model selection when constructing a dynamic momentum strategy may be less about maximizing returns and more about tailoring the strategy to align with individual risk preferences.

These results indicate that only marginal gains in risk-adjusted returns, if any, can be made by swapping a relatively well-performing volatility model for a more accurate one. As is shown in this paper, the relative order in performance is not consequent across stock universes. As a robustness check, we also applied our analysis to a momentum strategy with a 6-month sorting and holding period. The results, presented in Table 5 of the appendix, further support our main findings. Notably, the relative performance ranking of volatility models is not consistent across stock universes, reinforcing the conclusion that more accurate models yield, at best, only marginal improvements in risk-adjusted returns.

This implies that opting for a simple model, such as the realized variance, when implementing a dynamic momentum strategy yields similar returns while decreasing the necessary modelling work. The dynamic approach evaluated in this paper, adapted from Daniel

and Moskowitz (2016), requires computationally expensive procedures due to the large dataset and rebalancing of the portfolio is done every month. When applying a GARCH-model, this means estimating model parameters every month. Further, a realized variance model does not require any assumptions about the data to hold true.

As our implementation relies on simple price returns that account only for capital gains and not for dividends or other distributions, the resulting portfolio performance is not directly comparable to that of studies such as Daniel and Moskowitz or Jegadeesh and Titman. The omission of dividend-adjusted returns limits our ability to draw conclusions based on absolute return levels or Sharpe ratios. Nonetheless, our findings are consistent with Daniel and Moskowitz in terms of the directional impact that the dynamic strategy has on momentum performance. Importantly, the aim of this paper is to conduct a comparative analysis of volatility models within a consistent and controlled framework. By maintaining internal consistency in our data and methodology, we provide a solid basis for assessing the relative performance of the models, even if absolute performance metrics are not directly aligned with prior literature.

We argue the most important finding of this paper is the degree of volatility control that this dynamic momentum approach offers to investors. By adjusting the lambda parameter, an investor can choose a desired volatility level and, according to our results, achieve that level with consistency. In our implementation, the lambda value is estimated on a rolling basis, using only past data to dynamically scale the momentum portfolio's volatility to match that of the market. This ensures a strictly ex-ante approach and avoids look-ahead bias, setting our methodology apart from that of Daniel and Moskowitz (2016), where lambda is calibrated using full-sample information. This methodological improvement enhances the realism of our results and more accurately reflects the conditions under which an investor would apply the strategy in practice.

This "tailor-made" alignment of volatility allows investors to better integrate the momentum strategy within a broader portfolio context, particularly when managing overall risk exposure or adhering to specific risk budgets. It also provides a practical mechanism for making the strategy more accessible to risk-sensitive investors, such as institutions or individuals with strict draw-down constraints. Ultimately, the ability to precisely modulate volatility without compromising the ex-ante integrity of the strategy enhances both the usability and robustness of momentum investing in real-world applications.[2]

[2] We present a detailed discussion about the behavior of Lambda in the Appendix.

8 Conclusion

An examination of volatility forecasting accuracy and dynamic momentum strategy performance reveals that the most accurate volatility model does not necessarily deliver the highest risk-adjusted returns. This motivates considerations outside of forecasting accuracy alone when performing model selection for a dynamic momentum strategy.

This paper strengthens the case that volatility modelling can substantially improve momentum strategy performance, in line with earlier research by Daniel and Moskowitz (2016) and Barroso and Santa-Clara (2015). Our results show that while dynamic weighting based on volatility forecasts improves Sharpe ratios across both US and UK markets, the specific choice of volatility model yields only marginal differences in performance. This suggests that the benefits of dynamic momentum are robust across a range of models, and that relatively simple volatility specifications can be sufficient, an insight with practical value for both researchers and practitioners.

While this thesis offers valuable insights into the application of volatility forecasting within dynamic momentum strategies, it also opens several avenues for future research. One such direction involves examining the interaction between market regimes and the effectiveness of dynamic weighting. Future studies could analyze how model performance varies across bull, bear, and high-volatility regimes, and whether regime-specific adjustments to the lambda parameter could improve overall outcomes. Additionally, while this paper focuses on monthly rebalancing, more granular analyses using weekly or daily data may reveal nuances in how forecast timing affects portfolio behavior. Finally, evaluating whether certain models consistently perform better during crisis periods could offer practitioners a framework for model switching, tailored to real-time risk environments, potentially forming the basis for adaptive, volatility-aware investment strategies.

9 References

9.1 Papers

Barberis Nicholas, Schleifer Andrei, Vishny Robert W., 1997, A Model of Investor Sentiment, Working Paper 5926, National Bureau of Economic Research.

Barroso Pedro, Santa-Clara Pedro, 2014, Momentum has its moments, Journal of Financial Economics 116, 111-120.

Cooper, Gutierrez, Hameed, 2004, Market States and Momentum, The Journal of Finance, Vol. 59, No. 3.

Daniel Kent, Moskowitz Tobias J, 2016, Momentum crashes, Journal of Financial Economics 122, 221-247.

Gupta Tarun, Kelly Brian T, 2018, Factor Momentum Everywhere, Working Paper, Yale ICF No. 2018-23.

Hansen Peter R, Lunde Asger, 2005, A forecast comparison of volatility models: does anything beat a GARCH(1,1)?, Journal of Applied Econometrics, Vol.20, No. 7, 873-889.

Jegadeesh Narashiman, 1990, Evidence of Predictable Behavior of Security Returns, The Journal of Finance Vol. 45, No. 3, 881-898.

Jegadeesh Narashiman, Titman Sheridan, 1993, Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency, The Journal of Finance, Vol. 48, No. 1, 65-91.

Kelly Bryan T, Moskowitz Tobias J, Pruitt Seth, 2021, Understanding momentum and Reversal, Journal of Financial Economics 140, 726–743.

Lehman Bruce N, 1988, Fads, Martingales, and Market Efficiency, Working Paper 2533 National Bureau of Economic Research.

Moreira Alan, Muir Tyler, 2017, Volatility-Managed Portfolios, The Journal of Finance, Vol. 72, No. 4, 1611-1643.

Glosten Lawrence R, Jagannathan Ravi, Runkle David E, 1993, On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks, The Journal of Finance, Vol. 48, No. 5.

Scowcroft Alan, Sefton James, 2005, Understanding Momentum, Financial Analysts Journal, Vol. 61, No. 2.

9.2 Other

- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, CJ Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E.A. Quintero, Charles R Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors, 2020 SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python, Nature Methods, 17(3), 261-272. DOI: 10.1038/s41592-019-0686-2.

10 Appendix

10.1 Discussion of the Behavior of Lambda

The lambda discussed in this section refers to the one in the dynamic weight function from Daniel and Moskovitz (2016). The authors use it as follows: they fit an in-sample strategy to a specific lambda value that sets the strategy's volatility equal to that of the CRSP value-weighted index. After obtaining this value, they use it for out-of-sample estimates.

Our implementation differs slightly, as we found it problematic to fit the entire length of the value-weighted index, which could introduce lookahead bias. For example, if we want to make an optimal investment decision in 1980, we cannot use a lambda obtained from market data up to 2013, as that data would not exist yet. Therefore, we have synchronized the estimation of lambda with the estimation of the return regression and GARCH periods, always using a 5-year lookback period from any point in time.

Unlike in the model by Daniel and Moskovitz (2016), not only our weights vary over time, but the lambda used to define those weights does as well. Analyzing these results is foundational for further optimizing momentum strategies. First, we will discuss an alternative interpretation of lambda, then show how this understanding echoes in the way the algorithm fits the data, and the insights we can derive.

To inspire a quick intuition, we have prepared the following chart. All six subplots have the same axes and magnitudes. On the x-axis, we show the range of predicted daily returns obtained from the predictive return regression, spaced evenly into twenty ticks. On the y-axis, we show the range of predicted daily variances, obtained from the Realized Variance method[3], also spaced evenly into twenty ticks. The heatmaps themselves illustrate the weight that the dynamic weight function would assign to the risky asset (momentum portfolio) at any given predicted return - predicted variance pair. The individual heatmaps, arranged in a left-right, top-bottom orientation, demonstrate how changes in the lambda alone would affect our weight selection.

[3] The specific method is not crucial in this context; our focus is simply on obtaining a realistic daily variance range.

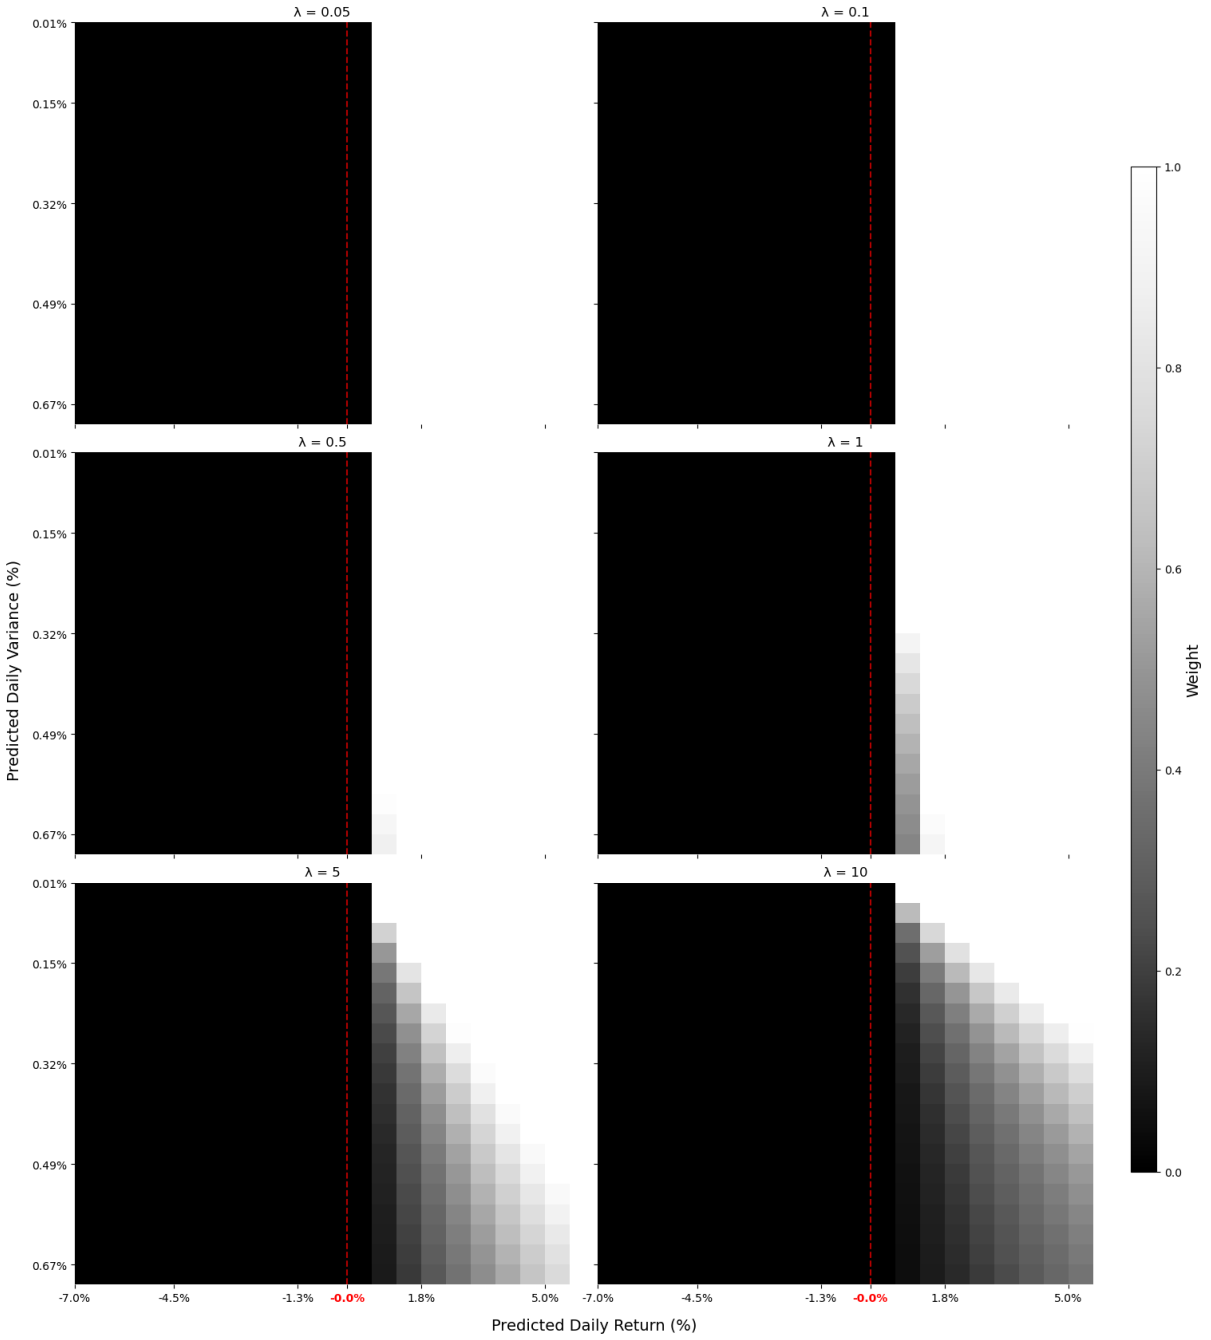


Figure 15. Effect of Lambda on Weight Selection. The chart illustrates how the dynamic weight function behaves as we vary lambda, holding all other factors constant. The x-axis represents a range of next-period return estimates, while the y-axis shows a range of next-period variance estimates. The weights on the risky asset (momentum portfolio) are constrained between 0 and 1. When return predictions are negative, no allocation to the risky asset is optimal. However, as return predictions turn positive, the function becomes more cautious as lambda increases, even when the underlying information remains unchanged.

We restrict our weights between 0 and 1—in other words, we do not allow for shorting the momentum strategy while taking credit on the risk-free asset. Therefore, the left side of zero on the predicted return axis results in a zero weight on the risky asset. If we allowed shorting, this side would show up as increasingly negative weights. If we have a good-quality negative

estimate for the next period's return, we want to hold off on investing anything. This finding is independent of the choice of lambda.

On the positive side with low lambda values, the dynamic weight function essentially becomes an on-off switch, recommending either a 0 or 1 weight on the risky asset. However, as we increase lambda by several orders of magnitude we observe drastic changes in behavior, as the dynamic weight function becomes more and more conservative.

The desired top-right regions (high predicted return coupled with low predicted variance) always remain at a weight of 1 on the momentum strategy. But the rest of the right side becomes more shaded, staying true to our preferences: we like higher returns but dislike high variance. Lambda acts as a control knob for the strategy. By changing lambda, we can influence the attitude of the strategy.

This finding only makes sense because our estimates of both return, and variance are imperfect. If they were perfect, there would be no reason to restrict investment—we would go all-in on the risky asset whenever we foresee a positive return. Since this is not the case, lambda allows for an elegant handling of uncertainty. If we are comfortable with our estimates, or have a higher risk appetite, we can reduce lambda, resulting in a more aggressive investment style—and vice versa.

The goal of this paper is to compare the effect of different volatility estimators on the momentum strategy introduced by Daniel and Moskovitz (2016); therefore, we did not explicitly model lambda conditional on market states. However, as it naturally emerged during the development of the codebase, we implemented a version where lambda was treated as a constant hyperparameter. Even this shallow implementation, using a higher-order lambda, delivered improvements.

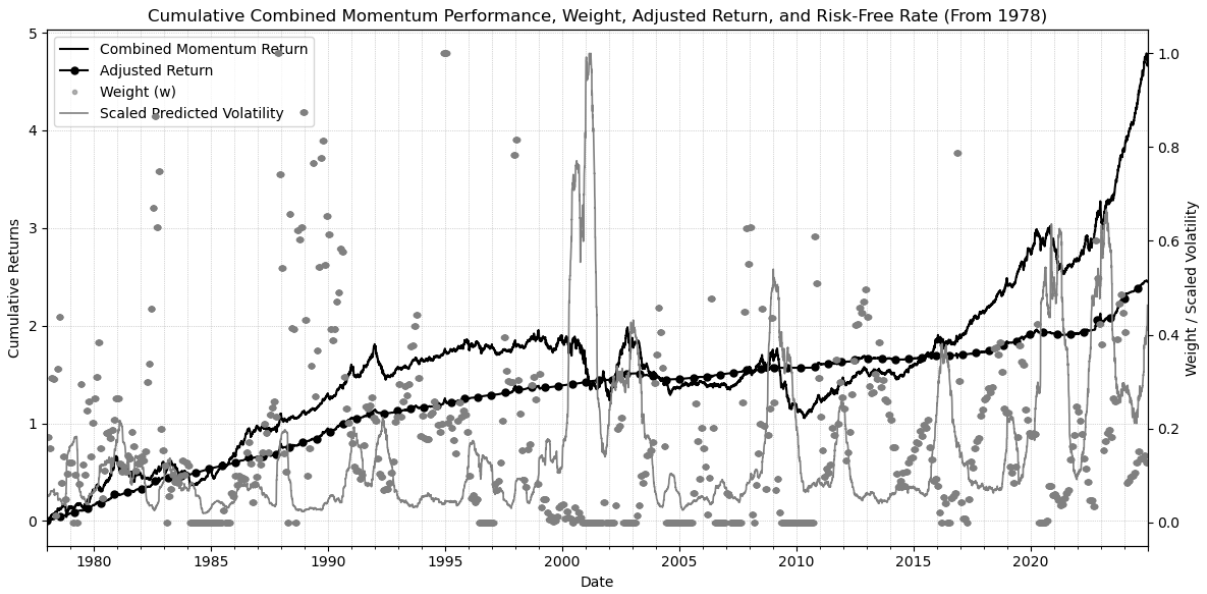


Figure 16. Dynamic Weight Allocation Strategy with Constant Lambda. By treating lambda as a constant hyperparameter, we obtain a smooth performance curve characterized by competitive returns and very low volatility (adjusted return series). In contrast, we also present the unadjusted momentum strategy. The model is highly responsive to market shocks, avoiding investment in the risky asset during crisis periods, as illustrated by the time series of weights. This results in a strategy resembling that of a high-return, low-risk bond. We present the results for the US market from 1978, with lambda set to 2.

Table 3. The numerical results comparing the performance of the constant lambda momentum strategy with the unadjusted strategy and the market capitalization-weighted market performance, spanning from 1978 to 2024.

Type	Annual Excess Return	Annual Standard Deviation	Sharpe Ratio
Constant Lambda Strategy	1.7%	2.4%	0.68
Momentum	4.7%	11%	0.43
Market	2%	17.2%	0.11

10.2 Trading Cost Considerations

One aspect we have not yet discussed, but which could be instrumental in explaining the persistence of momentum profits, is trading costs. This strategy requires trading huge volumes of stocks every month. Even though transaction costs have plummeted over the course of the 2000s, the cumulative impact remains substantial and casts a new light on discussions about profitability. Methodologically, we do not claim to break new ground here but simply wish to emphasize the importance of including trading costs in our calculations.



Figure 17. Monthly Trade Frequency in the Momentum Strategy. The number of trades executed each month to implement the momentum strategy on the US market (NYSE + NASDAQ) may differ between the two time series. This discrepancy arises because, in some cases, a stock remains continuously in the best or worst bucket for years. In such cases, we hold the position without executing a new trade, even though the bucket size remains the same.

The number of trades is steadily rising, which is a direct consequence of the market expansion. The peak is reached around 2000 and remains between 300 and 400 times 2 for the rest of the time series. Estimating trading costs for these trades is an inherently uncertain exercise due to the multitude of variables involved, many of which are difficult to quantify or predict.

The term trading cost is composed of many different components, and each of these components is likely influenced by geographically and legally dependent factors. For example, we would need to consider where we trade, who we trade with, what quantities are involved, and in some cases, even the stock's price, since some providers calculate trading costs differently for stocks priced above or below certain thresholds.

The components of the worst-performing portfolio would need to be sold short, which introduces additional costs and complexities. For example, in some EU countries, short selling by investment funds is not even permitted, meaning this strategy would be unfeasible for institutional investors.

Incorporating all this information and breaking it down by geography and governing law could easily warrant a separate paper on its own. Therefore, we decided to proceed with a simple approach: assuming a constant fee per trade, expressed as a percentage of the fund's return. We

call this variable trading_cost. With this assumption we just multiply the trading_cost by the number of trades and subtract it from the portfolio's cumulative log return at the beginning of each month.

This method has the undesirable side effect of introducing a large negative return on the first day of each month, which then propagates into the volatility calculation and, subsequently, the Sharpe ratio. In reality, a fund would incorporate trading costs into a reduced buy price and a later adjusted sell price when calculating returns, thus avoiding the adverse effects of a one-time trading cost deduction.

To address this, when calculating the Sharpe ratio, we use the return series that includes trading costs. However, the annual volatility is still calculated based on the original, unadjusted strategy time series. This approach ensures that the impact of trading costs is not overestimated.

We refer to the trading cost estimated above as the full trading cost. Since our entire approach employs dynamic weighting between the momentum strategy and the risk-free asset, we are typically not fully invested in the momentum strategy. Therefore, applying the full trading cost directly would significantly overestimate the impact.

To address this, we use the full trading cost time series and scale it by the corresponding weights, thereby trading costs change proportionally to the invested quantity. This assumption is certainly crude, but not entirely unrealistic.

By choosing the trading_cost variable, our approach is to incrementally increase its value and observe when the strategy collapses. If this breakdown occurs only at an unrealistically high level, the strategy is likely to hold up under real-world conditions, since actual trading costs would be far lower.

However, this is not the case. The strategy already collapses with the trading_cost variable set to 0.00001, or 0.001% per trade, which, from 2000 onward, results in an average 0.8% monthly reduction in portfolio value due to trading costs alone—*not even accounting for shorting costs yet*.

This means that for every $100 of portfolio value, we would pay 0.1 cent per trade and 80 cents per month for the full rebalance procedure on average. This is certainly not an unrealistically high estimate of transaction costs for a stock portfolio.

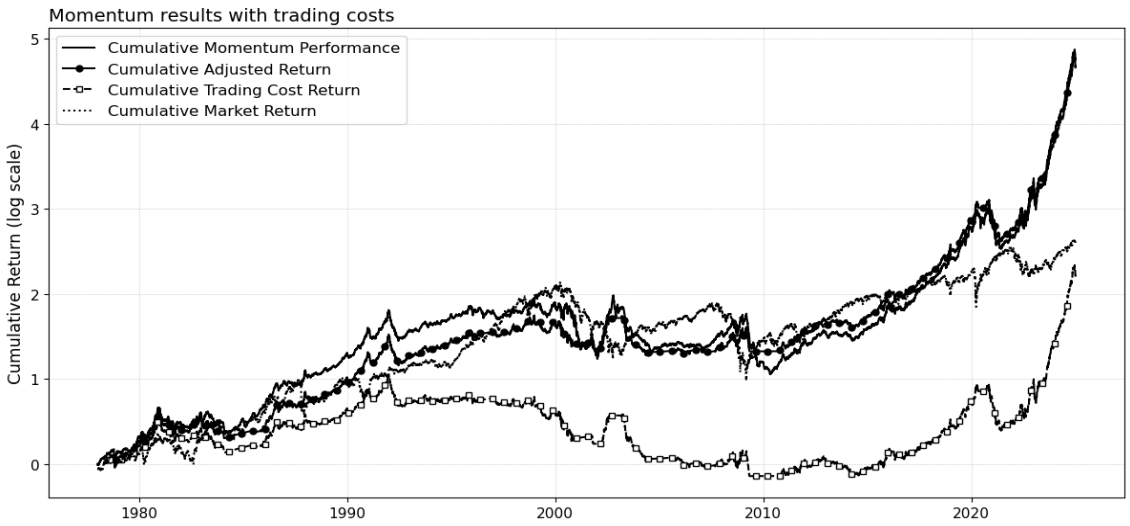


Figure 18. Comparison of US Momentum Portfolio Performance. The chart compares the performance of the US momentum portfolio from 1978, showing results for a strategy that includes trading costs versus one that does not, along with the performance of the US market and the unadjusted momentum performance over the same period.

Table 4. Numerical Comparison of US Momentum Portfolio Performance. This table presents a comparison of the performance of the US momentum portfolio from 1978, contrasting results for a strategy that includes trading costs with one that does not, as well as the performance of the US market over the same period.

Type	Annual Excess Return	Annual Standard Deviation	Sharpe Ratio
Momentum with Trading Costs	1.67%	9,57%	0.17
Momentum	4.7%	11%	0.43
Market	2%	17.2%	0.11

Even with the uncharacteristically strong performance of the past few years, we have only slightly outperformed a simple buy-and-hold, long-only market strategy in terms of the Sharpe ratio. However, we remain far behind the results achieved in earlier estimates. These outcomes are highly unstable; without the strong performance of the last few years, and with more realistic cost assumptions—including shorting costs and additional legal liabilities—the picture would look significantly worse.

In conclusion, momentum can be regarded as an artifact in the data. It certainly exists and, in theory, could be exploited or even further optimized. However, in practice, market frictions render it either impossible or, at the very least, economically unviable, potentially explaining why momentum profits continue to persist.

10.3 Six-month momentum portfolio results

Table 5.
Volatility forecast accuracy and strategy performance for a six-month holding period momentum portfolio.
These tables present forecast accuracy measures (RMSE and MAE) and portfolio performance metrics for the static momentum strategy and six dynamic strategies. Forecast accuracy metrics are computed in daily return units, while excess returns, volatility, and Sharpe ratios are annualized. Volatility predictions are scaled to match realized volatility. Panel A presents results for the US market (Nasdaq and NYSE) and Panel B presents results for the London Stock Exchange. Note: Return measures are based on price data only and do not include dividends. This may affect comparability with studies using total return data but does not limit the internal validity or purpose of this analysis.

Panel A.

US market.

Model	RMSE	MAE	Annual Excess Return	Annual Volatility	Sharpe-Ratio
Static Momentum	n/a	n/a	13.48%	11.59%	1.1632
Realized Variance	0.001720	0.001042	12.45%	10.60%	1.1748
GARCH	0.000180	0.000049	12.45%	10.60%	1.1743
EGARCH	0.002159	0.000987	12.51%	10.59%	1.1808
GJR-GARCH	0.000178	0.000048	12.45%	10.60%	1.1748
X-GARCH	0.000421	0.000244	12.47%	10.60%	1.1760
ARIMA-GARCH	0.011199	0.008725	12.38%	10.42%	1.1882

Panel B

London Stock Exchange.

Model	RMSE	MAE	Annual Excess Return	Annual Volatility	Sharpe Ratio
Static Momentum	n/a	n/a	23.40%	8.74%	2.6766
Realized Variance	0.001037	0.000656	21.66%	7.45%	2.9086
GARCH	0.000152	0.000030	22.26%	7.71%	2.9402
EGARCH	0.007076	0.003401	19.55%	7.42%	2.6357
GJR-GARCH	0.000121	0.000029	22.89%	7.72%	2.9654
X-GARCH	0.000374	0.000213	21.33%	7.78%	2.7420
ARIMA-GARCH	0.011101	0.008485	22.88%	7.69%	2.9737

10.4 Price Series Cleaning Procedure for Delisted Securities

```
def replace_repeated_prices(arraylike):
    counter = -3
    for _ in range(len(arraylike) - 3):
        if arraylike.iloc[-2] != arraylike.iloc[-1]:
            break
        elif arraylike.iloc[counter] == arraylike.iloc[counter + 1]:
            counter -= 1
        else:
            arraylike.iloc[(counter + 1) :] = np.nan
            break
    return arraylike
```

10.5 Artificial Intelligence Disclosure

During the preparation of this paper, we used generative Artificial Intelligence (AI) tools in several ways, as detailed below. In the research and discovery phase, we employed AI tools to summarize and pre-screen academic papers before reading them in full. This approach helped us filter out literature that was unrelated to our topic. All papers ultimately included in our research were thoroughly read by both of us. Using AI in this manner significantly increased the efficiency and breadth of our literature review, allowing us to stay focused and maintain productivity for other tasks. We did not identify any major risks associated with this use of AI.

All data collection was conducted manually on the Stockholm School of Economics campus using local computing resources. No AI tools were involved at this stage. We accessed the Refinitiv Eikon Database through the Refinitiv Eikon Microsoft Excel API and resolved any technical issues by consulting the documentation or online searches.

We developed the codebase using Visual Studio Code (VSCode), a free integrated development environment (IDE) provided by Microsoft. This software includes access to Microsoft Copilot, a subscription-based generative AI coding assistant. Users can select from a range of well-known large language model engines, such as ChatGPT, Gemini, or Claude. We primarily used Copilot for generating boilerplate code, particularly for producing charts and tables. In these cases, the immediate feedback provided by the tool allowed us to verify correctness in real time, which minimized risk. Given that creating professional-grade charts in Python can require between 50 and 100 lines of code, the ability to generate such scripts with a short prompt represents a significant productivity boost.

On a few occasions, we used Copilot to generate parts of the result-producing code. However, after encountering semantic errors, we limited this use. We estimate that over 90% of the code responsible for generating results was written by us or based on open-source packages. We conducted several code review sessions and thoroughly tested all result-generating code using debugging tools and, in some cases, unit tests. In our view, because generative AI currently lacks the ability to explain its outputs, the risk of semantic errors is especially high, particularly for inexperienced coders. As such, we approached AI-assisted coding with a high degree of caution.

The final area where we applied AI was in refining the grammar and flow of the text. We provided the AI tool with a neutral prompt instructing it not to alter the ideas or arguments presented, but rather to improve grammar, reduce repetition, and simplify unnecessarily complex sentences. If these improvements risked changing the intended meaning, the tool was directed to leave the text unchanged. This process enhanced the clarity and readability of the text with minimal risk, as all core ideas were originally written by us and we confirmed their integrity after revision.

We do not cite any AI tools, as we believe their role in this paper was purely technical and did not contribute substantively to the content.