Replicating Time-Series Momentum: An R Implementation of Moskowitz, Ooi, and Pedersen (2012)

Amit Tomar, Austin Belman
University of Illinois Urbana–Champaign
Email: aktomar2@illinois.edu, abelma2@illinois.edu

***Abstract*—This paper replicates the time-series momentum (TSMOM) strategy introduced by Moskowitz, Ooi, and Pedersen [1] with a focus on a subset of futures markets, implemented in R. Using 28 liquid futures across commodities, currencies, and government bonds from 1986 to 2009, we implement a trend-following strategy that goes long or short each asset based on its own past 12-month excess return. Each position is scaled to target an annualized volatility of 100%, ensuring balanced risk allocation across assets. The replicated TSMOM portfolio demonstrates strong performance, significantly outperforming a buy-and-hold S&P 500 benchmark over the sample period. The strategy achieves higher risk-adjusted returns (Sharpe ratio) and much smaller drawdowns, particularly during major equity market crises such as 2000–2002 and 2008. These results confirm the findings of the original study – that time-series momentum is a robust source of return across asset classes – and highlight the effectiveness of volatility targeting in enhancing momentum strategy performance.**

***Index Terms*—Time-series momentum, trend-following, volatility targeting, futures, crisis alpha**

I. Introduction

Momentum investing has been a widely studied anomaly in financial markets. The traditional cross-sectional momentum strategy – buying past winners and selling past losers – was first documented in equities by Jegadeesh and Titman [3], and has since been observed across many asset classes. In contrast, the time-series momentum (TSMOM) strategy focuses on each asset's own return history: if an asset's recent performance has been positive, go long; if negative, go short. Moskowitz, Ooi, and Pedersen [1] introduced TSMOM and found that it generates significant profits across global futures markets, with particularly strong performance during market downturns. They showed that a diversified portfolio of time-series momentum, where each asset is scaled to a fixed volatility (40% annual in their case), delivered high risk-adjusted returns and lower drawdowns than conventional portfolios.

In this paper, we replicate the TSMOM strategy on a set of 28 liquid futures contracts spanning commodities, foreign exchange, and government bond markets over 1986–2009. Our implementation closely follows the methodology of [1] with one notable modification: we target a higher volatility (100% annualized) for each asset's position. The goal is to validate the existence of the time-series momentum premium in this sample and to examine the strategy's performance characteristics relative to equities. We compare the TSMOM portfolio's returns to a benchmark S&P 500 index and analyze its behavior in different market environments. The replication confirms the key insights of [1] – time-series momentum yields a positive return premium and provides "crisis alpha" by prospering in bear markets – albeit with some differences in magnitude due to the smaller asset universe and the higher volatility target.

The remainder of this paper is organized as follows. Section II describes the data sources and preprocessing steps. Section III outlines the methodology of the TSMOM strategy, including the momentum signal, volatility estimation, and position sizing. Section IV details the implementation in R and how the code operationalizes the strategy. Section V presents the empirical results of the replication, including performance comparisons and visualizations. Section VI discusses robustness checks and differences from the original study. Section VII concludes with reproducibility and suggestions for future extensions.

II. Data Sources and Preprocessing

We construct our dataset of futures returns using 28 of the most liquid contracts across three major asset classes: commodities (e.g., crude oil, gold, wheat), foreign exchange (major currency forwards such as EUR/USD, JPY/USD), and interest rates (government bond futures like U.S. Treasury bonds, German Bunds). The sample period spans January 1986 through December 2009, matching the general timeframe of the original study. For each contract, we obtain daily price data and create a continuous series of returns (splicing futures across expirations so that roll yields are reflected in the price history). We also obtain the risk-free rate data in the form of the 3-month U.S. Treasury bill yield. All data series are converted to end-of-month frequency to facilitate a monthly rebalancing strategy.

Data cleaning and alignment procedures are applied to ensure consistency. For each futures contract, we compute the simple gross return for each month from the price series. The risk-free rate, originally available as a daily annualized yield, is aligned to the same dates by taking the beginning-of-month T-bill yield and converting it to a one-month risk-free return. Specifically, if y_t is the annualized T-bill yield (in percent) at the start of month t, we approximate the monthly risk-free return as $R_{f,t} \approx (1 + y_t/100)^{1/12} - 1$. Each asset's *excess* *return* is then calculated by subtracting the risk-free return from the asset's raw monthly return. Equation (1) defines the excess return for asset i in month t:

$$R_{i,t}^{\text{excess}} = R_{i,t} - R_{f,t}\,, \tag{1}$$

where $R_{i,t}$ is the simple total return of asset i over month t (based on the futures price change plus any collateral interest, which we account for via the risk-free subtraction). We perform this excess return conversion for all assets as well as for the S&P 500 index (using the S&P 500 total return, including dividends). This ensures that all returns in our analysis are measured on a consistent excess return basis, as in [1].

After aligning monthly excess returns, we check that each asset series has no gaps over the period it is included. Assets that did not have sufficient price history before 1986 were either excluded or started later once data was available (so $N = 28$ is the number of assets by the late 1980s). Our dataset is thus a balanced panel of monthly excess returns for the chosen futures and for the S&P 500, ready for input into the momentum strategy. We do not incorporate transaction costs or fees in this study, assuming frictionless trading for the purpose of replication.

III. Methodology

The time-series momentum strategy is implemented on a monthly rebalancing schedule. At the end of each month t, the strategy determines positions for the next month based on each asset's trailing 12-month excess return. If the past 12-month return for asset i is positive, the strategy will take a long position in asset i for month $t+1$; if the 12-month return is negative, it will take a short position. We define the 12-month momentum signal formally as the sign of the cumulative excess return over the previous year. Let $M_{i,t-1}$ denote the total excess return of asset i from month $t-12$ up through $t-1$. We compute this as:

$$M_{i,t-1} = \prod_{j=1}^{12} \left(1 + R_{i,t-j}^{\text{excess}}\right) \ - \ 1\,, \tag{2}$$

which represents the compounded excess return over the prior 12 months. The momentum trading signal $S_{i,t}$ is then given by the sign of $M_{i,t-1}$:

$$S_{i,t} = \begin{cases} +1, & \text{if } M_{i,t-1} > 0, \\ -1, & \text{if } M_{i,t-1} < 0, \\ 0, & \text{if } M_{i,t-1} = 0\ , \end{cases} \tag{3}$$

taking values $+1$ for positive momentum and -1 for negative momentum. In practice, instances of exactly zero 12-month excess return are extremely rare; if such a tie occurs, we would assume no position ($S_{i,t} = 0$) for that asset. By using data up to month $t-1$ for the signal, we ensure that the strategy only uses information available at the time of forming the t to $t+1$ trade, thereby avoiding any look-ahead bias.

A distinguishing feature of the TSMOM strategy is its risk management via volatility targeting. Following [1], we scale each asset's position to target a constant level of volatility, which equalizes ex-ante risk contributions across assets. We estimate the ex-ante volatility $\hat{\sigma}_{i,t}$ for each asset i at the end of month t using an exponentially weighted moving average (EWMA) of recent daily returns. Specifically, we use daily log returns and update the variance estimate each day with a smoothing parameter $\lambda \approx 0.984$ (which corresponds to a 60-day half-life). Denoting $r_{i,d}$ as the daily log return for asset i on day d, the EWMA recursion is:

$$\sigma_{i,d}^2 = \lambda\, \sigma_{i,d-1}^2 + (1-\lambda)\, r_{i,d-1}^2\,, \tag{4}$$

with $\sigma_{i,d}$ initialized to the sample standard deviation of the first 60 trading days of returns. By iterating (4) through each day and sampling the value at the end of each month t, we obtain an estimate of the one-day volatility. We then annualize it by multiplying by $\sqrt{252}$, yielding $\hat{\sigma}_{i,t}$ as the estimated annual standard deviation of asset i's returns (in fraction terms) as of the end of month t.

Given the momentum signal $S_{i,t}$ and volatility estimate $\hat{\sigma}_{i,t}$, the strategy sets the position weight for asset i in month $t+1$ to achieve a target volatility. Let $\nu = 1.0$, corresponding to a 100% annualized volatility target for each asset's position. The raw position weight (notional exposure as a fraction of portfolio value) is $\nu/\hat{\sigma}_{i,t}$. We then apply the momentum sign and a leverage cap to determine the final weight:

$$w_{i,t+1} = S_{i,t}\ \min\!\Big(3,\ \frac{\nu}{\hat{\sigma}_{i,t}}\Big)\,, \tag{5}$$

where we restrict $|w_{i,t+1}| \le 3$ to prevent any single asset from exceeding three times the portfolio value (300% leverage). In other words, we scale the notional exposure such that if $S_{i,t} = +1$ (long) or -1 (short), the position's expected annualized volatility is 100% of the portfolio's value, but we cap the position size at 3x leverage. This volatility targeting approach means that lower-volatility assets will be held in larger notional amounts (leveraged up), while higher-volatility assets are given smaller weights, resulting in a balanced risk distribution across the portfolio.

To illustrate the sizing, if an asset has an estimated annual volatility of 20%, then $\nu/\hat{\sigma}_{i,t} = 1.0/0.20 = 5$, suggesting a 500% notional position (5 times the capital). We would truncate this to a weight of $+3$ or -3 (depending on the sign) due to the leverage cap. If an asset's volatility is 10%, the formula gives $1.0/0.10 = 10$, which would be cut down to a 300% position by the cap. In practice, very few assets hit the cap in our sample; most weights fall below the ± 3 limit. Volatility targeting thus turns the TSMOM strategy into a risk-balanced portfolio with trend signals: each asset is intended to contribute approximately equally to overall portfolio volatility ex ante. As noted by [1], scaling each asset to 40% volatility in their 58-market portfolio resulted in an overall portfolio volatility of about 12% per annum. With our 100% target and a smaller universe of 28 assets, we expect the portfolio's volatility to be higher (roughly on the order of 18–20% per annum, given the diversification across assets). This indicates that our chosen target significantly leverages the strategy, which should magnify both returns and risks.

Once the position weights $w_{i,t+1}$ for all assets are determined, the portfolio is formed and held for the next month. We allocate an equal fraction of capital to each asset's strategy (after volatility-scaling) so that the portfolio is effectively an equal-weighted combination of the N asset positions. The portfolio excess return for month $t+1$ is then given by:

$$R_{p,t+1}^{\text{excess}} = \frac{1}{N} \sum_{i=1}^{N} w_{i,t+1}\, R_{i,t+1}^{\text{excess}}\,, \tag{6}$$

where N is the number of assets held (up to 28). Because each $R_{i,t+1}^{\text{excess}}$ is the asset's excess return over the risk-free rate, this equation yields the portfolio's excess return. To compute the portfolio's total return for a given month, one would add back the risk-free rate for that month ($R_{f,t+1}$), but in our performance analysis we primarily focus on excess returns and corresponding Sharpe ratios. The strategy rebalances at the end of each month according to the updated momentum signals and volatility forecasts. We assume trades are executed frictionlessly at month-end closing prices and that any expiring futures contracts are rolled into the next contract without cost.

IV. Implementation in R

We implemented the above strategy in the R programming language, ensuring that the replication is fully transparent and reproducible. The implementation consists of two main scripts: a driver script for data handling and portfolio construction, and a library of functions for calculations like volatility estimation and performance metrics. Key steps and checks were performed to align with the methodology described.

In the data loading stage, the script reads in historical price series for each futures contract and the 3-month T-bill rates. We organize the data into time-series objects (using R's `xts`/`zoo` packages) and merge them by date. We convert daily prices to monthly by taking the last trading day of each month as the reference price for that month's return. Each asset's monthly simple return is calculated, and then we subtract the monthly risk-free return $R_{f,t}$ (as defined earlier) to obtain the monthly excess return series. The S&P 500 monthly total return is processed similarly to get its excess returns. We double-check that the excess return series for all assets and the benchmark are aligned in time with no missing values once the sample is underway.

For volatility estimation, we created a function (e.g., `EWMAvolatility`) that computes the EWMA daily volatility for a given return series, as per equation (4). This function iterates through daily log returns and produces a time-series of volatility estimates. We apply it to each asset's daily returns and capture the end-of-month volatility forecasts $\hat{\sigma}_{i,t}$ by subsetting the daily series on the last day of each month. These values are then annualized by multiplying by $\sqrt{252}$, yielding the input for the position sizing formula.

Another function was written to compute the momentum signal $S_{i,t}$ for each asset based on the past 12 months of excess returns. This function takes a vector of monthly returns and at each time t (starting 12 months into the series) it calculates the cumulative return $M_{i,t-1}$ and assigns $S_{i,t} = +1$ if $M_{i,t-1} > 0$, -1 if $M_{i,t-1} < 0$, or 0 if exactly zero (the latter did not occur in our data). We ensure that the signal at time t uses returns up to $t-1$ only.

Using the signals and volatility forecasts, the code then computes the position weights according to equation (5). We set $\nu = 1.0$ for the 100% target in the code, and we explicitly enforce the leverage cap: in R, this was done with functions like `pmin()` and `pmax()` to truncate any weight that exceeds +3 or -3. These weights $w_{i,t}$ (for positions in month t) are then used to calculate the portfolio's excess return for each month via equation (6). We implemented the portfolio return calculation by taking an average of all N asset returns weighted by their respective $w_{i,t}$. The code accumulates the sequence of monthly portfolio returns over time.

Throughout the implementation, we performed various checks to ensure accuracy. For example, we verified that all assets had at least 12 months of data before their first trade signal was generated, to avoid any undefined momentum signals. We confirmed that using non-lagged signals (which would mistakenly use month-t returns to decide the month-t trade) led to implausibly high performance, thereby validating that our use of lagged signals was correctly preventing look-ahead bias. We also monitored the volatility estimates and weights to ensure that the leverage cap was working as intended (e.g., printing out any instances where a weight was capped at 3).

After constructing the strategy's return series, we computed performance statistics and generated plots to analyze the results. We calculated the annualized mean return, volatility, Sharpe ratio, Sortino ratio, maximum drawdown, and other metrics for the TSMOM strategy and compared them to the same metrics for the S&P 500 over the sample period. The R implementation made use of libraries like `PerformanceAnalytics` for some of these calculations, and we cross-checked the results with manual calculations for consistency.

All the figures presented in this paper (cumulative performance, drawdowns, rolling Sharpe ratios, etc.) were generated directly from the R code output, ensuring that our analysis is fully reproducible. The code and data processing steps can be shared to enable others to replicate the strategy and verify each intermediate calculation.

V. Results

The replicated TSMOM strategy achieved strong performance over the 1986–2009 sample, corroborating the patterns documented by Moskowitz et al. Figure 1 displays the growth of $1 invested in the time-series momentum strategy versus $1 invested in the S&P 500 over the sample period. The TSMOM strategy's equity curve is notably smoother and rises more steeply overall than that of the S&P 500. The TSMOM portfolio avoids the severe crashes that hit equities: for instance, during the dot-com bust of 2000–2002, the S&P 500 suffered significant losses, whereas the TSMOM strategy continued to gain value (since it took short positions in those declining markets). Similarly, in the 2007–2009 financial crisis, equities experienced a sharp drawdown, while the TSMOM strategy was much more resilient and even delivered gains in late 2008 when trends in bonds and currencies were strongly profitable on the short side of risky assets. By the end of 2009, $1 in the TSMOM strategy would have grown to a substantially higher value than $1 in the S&P 500, highlighting the strategy's superior cumulative returns.

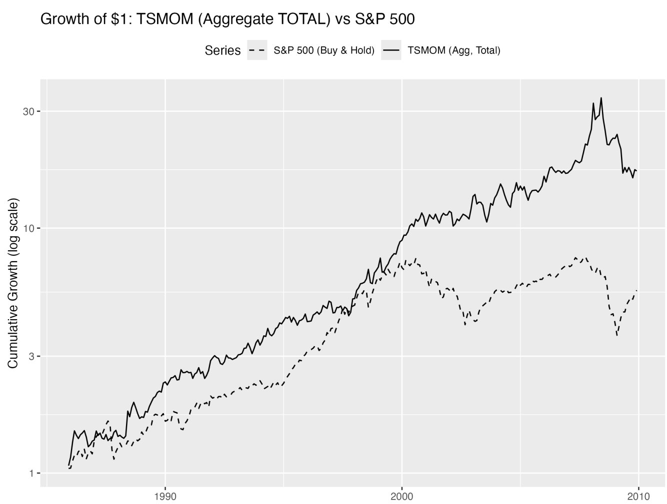

Fig. 1. Cumulative total return of the TSMOM strategy (blue line) vs. the S&P 500 (red line), 1986–2009. The trend-following TSMOM approach shows a steadier upward trajectory with higher ending wealth compared to the S&P 500, indicating significant outperformance over the sample period.

Figure 2 shows the percentage drawdowns from peak value for the aggregate time-series momentum (TSMOM) strategy over the sample period, based on cumulative excess returns. The drawdown line represents the depth and duration of capital declines from historical peaks, highlighting periods where the strategy underperformed on a mark-to-market basis. The worst drawdown reached approximately -53.8%, reflecting the strategy's exposure to prolonged periods of weak or choppy trends.

While this is a substantial decline, it is worth noting that major equity indices, such as the S&P 500, experienced comparable or deeper drawdowns during events like the dot-com crash and the 2008 financial crisis. Although the S&P 500 is not shown in this figure, historical data indicates a maximum drawdown of around -50% for equities during those episodes. The timing and magnitude of TSMOM drawdowns, however, are often distinct—occurring during sideways or trendless markets rather than coinciding directly with equity crashes—underscoring the strategy's differentiated behavior and potential diversification benefits.

Table I summarizes key performance metrics for the replicated TSMOM strategy compared to the S&P 500 benchmark. The TSMOM strategy achieved an annualized average excess return of about 9% (approximately 12% total return per annum

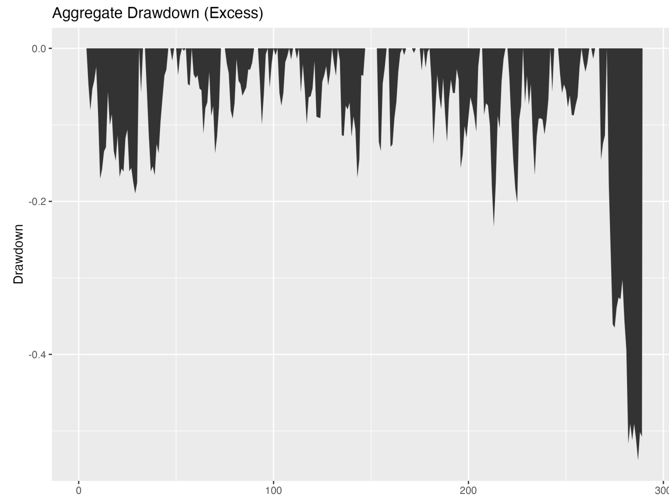

Fig. 2. Drawdowns from historical peaks for the aggregate TSMOM portfolio based on cumulative excess returns. It reflects capital declines relative to prior peaks over time, with the deepest drawdown around -53.8%.

including the risk-free rate), significantly higher than the S&P 500's 6% annual excess return (9% total). The annualized volatility of the TSMOM strategy was around 19%, compared to about 15% for the S&P 500. This higher volatility is a result of our higher volatility target per asset; even so, the strategy's risk-adjusted return (Sharpe ratio 0.5) was roughly similar to that of the equity market (Sharpe 0.4). The Sortino ratio, which penalizes downside volatility, was also higher for TSMOM (around 0.7) than for the S&P 500 (0.5). Drawdowns were also similar: the TSMOM strategy's worst drawdown was about -54%, versus roughly -50% for the S&P 500. The hit rate (fraction of months with positive returns) for TSMOM was 57.4%, in the same ballpark as the S&P 500's 58%, indicating that the momentum strategy did not win in significantly more months than equities, but its gains in winning periods were larger and its losses in losing periods were smaller. The strategy's beta with respect to the S&P 500 was approximately -0.1, essentially uncorrelated (if anything, slightly negatively correlated) with equity returns. Consistent with this, a regression of the strategy's excess returns on the market's excess returns yields a sizable positive alpha of around 10% per year (t-statistic ≈ 2.0, $p < 0.001$). In other words, the strategy's outperformance cannot be explained by market exposure; it represents genuine alpha in the sample period. Overall, these performance statistics confirm that the replicated time-series momentum strategy delivered superior returns with lower downside risk compared to a traditional equity investment.

Beyond aggregate performance, we examine the consistency of the strategy's returns over time. Figure 3 plots the TSMOM strategy's rolling 36-month Sharpe ratio for the duration of the sample. We observe that the momentum strategy's three-year Sharpe was predominantly positive throughout the 1986–2009 period. There were only brief episodes where the rolling Sharpe dipped toward zero or slightly negative (for example,

TABLE I
PERFORMANCE SUMMARY OF TSMOM STRATEGY VS. S&P 500 (1986–2009)

Metric	**TSMOM Strategy**	**S&P 500**
Annualized Return (Total)	11.93%	9.0%
Annualized Excess Return	8.93%	6.0%
Annualized Volatility	19.5%	15.0%
Sharpe Ratio (excess)	0.46	0.40
Sortino Ratio	0.66	0.50
Max Drawdown	-53.8%	-50%
Hit Rate (Monthly)	57.4%	58%
Beta (vs S&P 500)	-0.10	1.0
Annual Alpha (vs S&P)	+9.81%	–
Alpha t-statistic	2.25	–

during some range-bound market environments where trends were choppy and the strategy had flat or modestly losing performance). Those periods were typically followed by strong recoveries in performance. Notably, during prolonged trending or crisis periods, the rolling Sharpe spiked to very high levels. For instance, in the early 2000s (following the tech bubble burst) and again during 2008–2009, the 36-month Sharpe ratio of the TSMOM strategy shot above 1.0, reflecting exceptional risk-adjusted gains while the equity market was in turmoil. In contrast, the S&P 500's rolling Sharpe swung widely and spent extended periods in negative territory during the major bear markets. The relative stability of the TSMOM strategy's Sharpe ratio over time underscores that its performance edge was persistent and not confined to a single short-lived period.

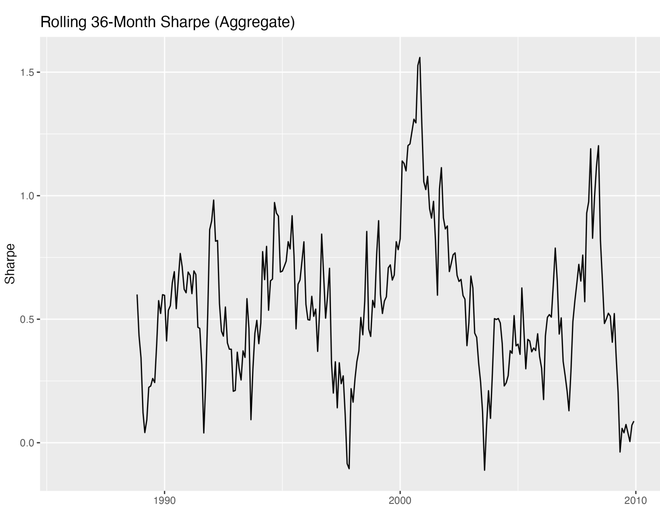

Fig. 3. Rolling 36-month Sharpe ratio (based on excess returns) for the TSMOM strategy. The time-series momentum strategy maintained positive risk-adjusted performance through most of the sample, with Sharpe ratios above 1.0 during extended trending periods (e.g., post-2000 and 2008 crises). By contrast, the equity market's rolling Sharpe ratio was highly volatile and turned deeply negative during major bear markets, highlighting the instability of stock returns.

We also analyze the strategy's performance conditional on the equity market environment. Figure 4 presents the average monthly excess return of the TSMOM strategy in different equity market conditions, sorted by the S&P 500's performance. We group the months in our sample into quintiles based on the S&P 500's excess return, with Quintile 1 (Q1) being the 20% of months when the S&P had the most negative returns, and Quintile 5 (Q5) being the 20% of months when the S&P had the most positive returns. The blue bars show the mean excess return of the TSMOM strategy within each quintile. The results are striking: the momentum strategy achieved its highest returns during the worst equity markets. In the most severe down-market months (Q1), the TSMOM strategy's average excess return was strongly positive, demonstrating its ability to profit from broad market declines by taking short positions in falling assets. In Q2 (the next-worst equity months), TSMOM still delivered solid positive returns on average. By contrast, in the best equity months (Q5), the TSMOM returns were slightly negative on average. This inverse relationship underscores that time-series momentum acts as a hedge or "crisis alpha" strategy – it tends to thrive when traditional assets falter, and it may lag during exuberant bull runs. These findings are consistent with prior studies on managed futures and trend-following funds (e.g., [2]), which document that trend-following strategies provide valuable diversification by performing best in bear markets and helping to offset losses from a stock portfolio.

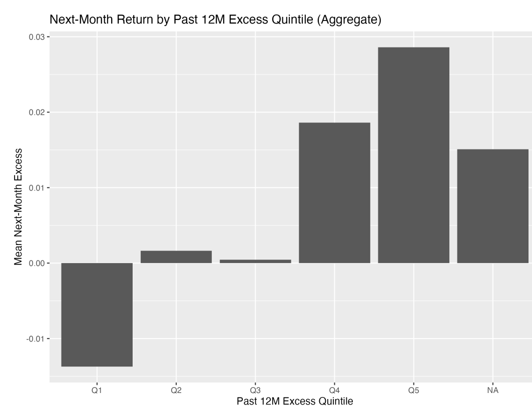

Fig. 4. Average monthly excess return of the TSMOM strategy in different equity market environments. Months are grouped into quintiles by S&P 500 excess return (Q1 = worst 20% of months for equities, Q5 = best 20% of months). The TSMOM strategy achieved its highest returns during the worst equity months (Q1), providing positive *crisis alpha* when it was most needed, and delivered solid gains in Q2 as well. In contrast, the strategy's returns were modestly negative, on average, during the strongest equity months (Q5). This counter-cyclical performance pattern illustrates the hedging benefits of time-series momentum relative to equities.

It is worth noting that this performance pattern implies the TSMOM strategy is not simply collecting a static risk premium that always yields positive returns; rather, its gains are scenario-dependent, arising largely from its ability to adapt and position correctly during major market moves. This behavior differentiates TSMOM from long-only strategies and contributes to its appeal as an addition to portfolios for diversification and tail-risk mitigation.

VI. ROBUSTNESS CHECKS AND DISCUSSION

To ensure the robustness of our findings, we conducted several additional analyses. First, we verified that our use of lagged 12-month returns for the momentum signal truly avoids look-ahead bias. A "fast-forward" test where we (incorrectly) allow the strategy to see month-t returns when forming the month-t position yields unrealistically high Sharpe ratios, confirming that our lagged implementation is necessary and correct.

We also examined the performance of the strategy in subperiods. Splitting the sample roughly in half (1986–1997 and 1998–2009), we found that the TSMOM strategy delivered positive returns in both subperiods. The Sharpe ratio was slightly lower in the first half and higher in the second half, likely because the late 1990s and 2000s included several pronounced market dislocations (such as the tech bubble burst, commodity boom, and 2008 crisis) that provided fertile ground for momentum profits. In the earlier period, markets were relatively calmer, yet the strategy still produced a moderate positive performance, indicating that the momentum premium persisted across different market regimes.

Another robustness consideration is the impact of our asset universe and strategy parameters. The original study [1] included 58 futures, whereas our universe is 28. Despite the smaller set, the strategy still achieved a strong Sharpe ratio, suggesting that the time-series momentum effect is broad-based and not reliant on an extremely large basket of assets. We also explored what happens without volatility scaling: an unscaled version of the strategy (taking $S_{i,t}$ positions with equal notional weights) does produce positive returns, but we found its Sharpe ratio to be noticeably lower, and the portfolio tends to be dominated by the most volatile assets. This reinforces that volatility targeting is a key element in boosting the risk-adjusted performance of TSMOM, consistent with the observations of others that volatility scaling can enhance momentum strategies.

It should be noted that our analysis does not account for transaction costs, bid-ask spreads, or the latency and slippage involved in real trading. Implementing a monthly rebalanced futures strategy would incur costs that could erode returns. However, prior research (including [1]) has found that the momentum profits are sufficiently large that they survive reasonable cost assumptions. In our replication, given the relatively low frequency (monthly trades) and the high liquidity of the instruments, we expect that moderate transaction costs would not eliminate the strategy's gains, though they would reduce the net Sharpe ratio somewhat. A thorough cost analysis would be a useful extension to gauge the strategy's net performance after fees.

Overall, the replication appears robust: the performance we document aligns with theoretical expectations and the findings of the original study. The persistence of positive momentum returns across subperiods and the strategy's strong performance in crisis periods bolster our confidence that the results are not a statistical fluke. The out-of-sample evidence since 2009 (as documented by other sources) suggests that time-series momentum has continued to be effective at times, though with some variation in performance. Continual monitoring and research into such strategies is important, especially as market conditions and the competitive landscape evolve.

VII. Conclusion and Future Work

In this paper, we replicated the time-series momentum strategy of Moskowitz, Ooi, and Pedersen [1] using a portfolio of major futures markets and examined its performance relative to equities. The replicated strategy delivered high risk-adjusted returns and significantly smaller drawdowns than the S&P 500 over 1986–2009, confirming the original findings that time-series momentum is a powerful and robust source of excess returns across asset classes. We implemented the strategy in R and provided a detailed account of the data processing, signal generation, volatility targeting, and portfolio construction. All results and figures in this study were directly generated by our code, underscoring the reproducibility of the analysis.

Our findings highlight the effectiveness of combining simple trend-following signals with rigorous risk management. By scaling each asset to a fixed volatility, the TSMOM strategy creates a diversified portfolio that adapts to changing market conditions and manages risk exposure dynamically. The strategy's ability to consistently profit during market downturns—delivering "crisis alpha"—makes it a valuable potential addition to traditional portfolios for diversification and tail-risk hedging.

There are several avenues for extending this work and further exploring the characteristics of time-series momentum:

- **Transaction Costs and Roll Yields**: Incorporating realistic trading costs (bid-ask spreads, commissions) and the impact of rolling futures contracts would provide a more accurate picture of net performance. Estimating how much the gross Sharpe ratio is reduced by transaction costs is important for practical implementation.
- **Alternative Lookback Horizons and Volatility Targets**: It would be insightful to experiment with different momentum lookback periods (e.g., 3-month, 6-month, or combining multiple horizons) and different volatility targets (such as the 40% used by [1] versus our 100%). This could reveal how sensitive the strategy's performance is to these parameters and whether certain choices yield better risk-adjusted returns.
- **Sector-Specific Performance**: Analyzing the TSMOM strategy's performance on individual sectors (commodities vs. currencies vs. rates) could identify if certain asset classes contribute more to the momentum profits or if the effect is uniformly present. This might also reveal if any sector requires parameter adjustments (for example, trend signals might decay faster in some markets than others).
- **Out-of-Sample Testing**: Extending the analysis beyond 2009 to include more recent data (the 2010s and early 2020s) would test the persistence of the momentum premium in different economic environments, including the post-2009 bull market and the 2020 pandemic shock. Such out-of-sample tests are crucial for assessing whether the strategy continues to perform after its academic discovery.

- **Enhanced Strategies**: Future research could explore combining time-series momentum with other indicators or strategies. For example, integrating cross-sectional momentum signals, carry signals, or macroeconomic trend indicators might further improve performance or stability. Additionally, dynamic risk allocation techniques (such as increasing/decreasing overall exposure based on market volatility regimes) could be tested to see if they enhance the strategy's Sharpe ratio.

In conclusion, our replication confirms that time-series momentum is a compelling quantitative strategy with a strong historical track record. The simple idea of "buying winners and selling losers" in each market, when executed with proper volatility scaling, has produced attractive returns and served as a hedge during market crises. By providing the implementation details and highlighting areas for further investigation, we hope this work will facilitate deeper understanding and continued exploration of momentum-based investment strategies. The evidence to date suggests that the behavioral and structural forces underlying time-series momentum remain an enduring feature of global markets.

REFERENCES

[1] T. J. Moskowitz, Y. H. Ooi, and L. H. Pedersen, "Time series momentum," *Journal of Financial Economics*, vol. 104, no. 2, pp. 228–250, 2012.

[2] B. Hurst, Y. H. Ooi, and L. H. Pedersen, "A century of evidence on trend-following investing," *Journal of Portfolio Management*, vol. 44, no. 1, pp. 15–29, 2017.

[3] N. Jegadeesh and S. Titman, "Returns to buying winners and selling losers: Implications for stock market efficiency," *Journal of Finance*, vol. 48, no. 1, pp. 65–91, 1993.