

Nonlinear Time Series Momentum

Tobias J. Moskowitz*

Yale University & NBER

Riccardo Sabbatucci[†]

Stockholm School of Economics

Andrea Tamoni[‡]

University of Notre Dame

Björn Uhl[§]

Universität Hamburg

December 17, 2025

ABSTRACT

We document a persistent nonlinear relationship between price trends and risk-adjusted returns across markets and asset classes that is consistent with asset pricing theory. Nonlinearities in time series momentum are consistent with past returns reflecting information about conditional expected returns, in line with investors using conditioning information to form efficient portfolios. Machine learning techniques are useful in uncovering these relationships and yield economically and statistically significant out-of-sample improvements in time series momentum strategies.

Keywords: Time series momentum; return forecasts; nonlinear regression; machine learning

JEL classification: G17, C53, C58, G11

*Department of Finance, Yale University and AQR Capital. **Email:** tobias.moskowitz@yale.edu

[†]Department of Finance, Stockholm School of Economics. **Email:** riccardo.sabbatucci@hhs.se

[‡]Department of Finance, Mendoza College of Business, University of Notre Dame. **Email:** ata-moni@nd.edu

[§]Faculty of Business Administration, Universität Hamburg. **Email:** bjoern.uhl@uni-hamburg.de

1 Introduction

Time-series momentum (TSMOM) is a well-documented empirical phenomenon in financial markets that applies across many markets and asset classes ([Moskowitz et al. \(2012\)](#), [Zakamulin and Giner \(2020\)](#), [Hurst et al. \(2017\)](#), [Huang et al. \(2020\)](#)).¹ TSMOM captures the continuation of an asset's *own* past returns, as opposed to cross-sectional momentum (XSMOM) ([Jegadeesh and Titman \(1993\)](#), [Asness \(1994\)](#), [Kelly et al. \(2021\)](#)), which reflects the relative performance of assets. This phenomenon is one of the chief empirical anomalies highlighted in discussions of market efficiency. The popular trend-following and managed futures (CTA) strategies of many hedge funds are based on this research, commanding many hundreds of billions of dollars in assets under management. Better understanding this empirical regularity is interesting to both academia and practice.

Theories for why the past returns of an asset are related to its future expected return have received both risk-based and behavioral motivations. Behavioral theories provide a compelling framework. Indeed, the earliest behavioral asset pricing models ([Daniel et al. \(1998\)](#), [Barberis et al. \(1998\)](#), [Hong and Stein \(1999\)](#)) are actually about time-series not cross-sectional momentum and use various behavioral biases from psychology to generate momentum of a single asset. Rational-based stories search for underlying sources of risk that give rise to a momentum premium, where past returns are an indicator of exposure to some compensated risk ([Boudoukh et al., 2025](#)). Empirically, TSMOM seems to have negative exposure to traditional sources of risk, such as equity and interest rate risk as characterized by TSMOM's well-known and well-documented U-shape pattern with equity and bond returns ([Moskowitz et al. \(2012\)](#), [Hurst et al. \(2017\)](#)). This feature of TSMOM is not present for XSMOM and presents a challenge to traditional risk explanations.

We take a different theoretical approach to the study of TSMOM. We examine it from an investment perspective, and ask how might an investor back out conditioning information about an asset's expected return from past returns. For example, [Ferson and](#)

¹Recent research focuses on crowding ([Bollen et al., 2021](#); [Baltas, 2019](#)), in the differences between cross-sectional and time series momentum ([Kwon and Satchell, 2021](#)), equity factors ([Babu et al., 2020](#)), and the impact of leverage ([Forner et al., 2018](#)) and volatility scaling ([Kim and Tse, 2016](#)).

[Siegel \(2001\)](#) show that an investor seeking to maximize expected returns for a given level of risk, should use a nonlinear function of any expected return signal to maximize the in-sample unconditional Sharpe ratio of the portfolio. Specifically, the unconditional efficient portfolio will weight expected return signals with an “S”-function, where signals close to zero expected returns are weighted linearly and more extreme return signals are down weighted. The intuition is that for extreme signals, the optimal weight scales or “rolls” back, because the investor trades off return prediction with variance reduction by not wanting to be too exposed to the risky asset. [Daniel and Moskowitz \(2016\)](#) derive a similar weighting function for timing the XSMOM factor (i.e., the momentum of momentum). A similar intuition and result also obtains if there is estimation error in the expected return signal, which will be larger for more extreme signals. Applied to TSMOM, this optimal weighting scheme implies a natural form of nonlinear momentum.

We examine the theoretically-motivated nonlinear momentum (NL) scheme against other empirical implementations of TSMOM from both the literature and machine learning techniques. The literature has largely focused on simple binary functions of positive or negative momentum based on threshold rules ([Moskowitz et al., 2012](#)), or linear signal weights, which are atheoretical. If the theory is accurate, then the nonlinear weighting scheme consistent with the model framework should outperform these simple weighting schemes. It is also interesting to see if statistical techniques such as machine learning can uncover the same nonlinear weighting implied by the theory, or even improve upon it. We compare how our theoretically-generated weights do against a set of weights derived from sophisticated neural networks that try to maximize the out-of-sample Sharpe ratio of a TSMOM trading strategy.

First, we show that nonlinear momentum is pervasive across a variety of asset classes and across horizons and frequencies. We examine 8 equity index futures markets, 24 commodity futures, and 21 interest rate and currency futures contracts and find that a nonlinear TSMOM weighting consistent with theory outperforms other TSMOM weighting schemes in the literature. Across lookback periods ranging from 1-month to 12-months, and across frequencies (daily, weekly, monthly) we find that NLTSMOM outperforms other TSMOM weights consistently and significantly. We also show that these results

obtain when analyzing equity long-short factor momentum (e.g., among the Fama and French five factors) as in [Gupta and Kelly \(2019\)](#) and [Ehsani and Linnainmaa \(2022\)](#).

Second, a key feature of TSMOM is its convex return relation with equity (and bond) markets. [Moskowitz et al. \(2012\)](#) find that TSMOM strategies produce their greatest returns in extreme markets, particularly low return environments, which provide attractive hedging properties to an investor holding the market portfolio or a 60-40 traditional stock-bond portfolio. We find that the NL weighting of TSMOM improves their hedging property as well, generating even greater outperformance in extreme down markets, which explains the majority of the improvement in unconditional mean returns from the NL weighting. This finding may be consistent with the intuition that motivates the NL weighting theoretically – it is designed to mute exposure to more extreme signals, which reduces the risk and noise associated with extreme signals and thus improves profitability. Consistent with this motive, the Sharpe ratio improvement is even larger than the improvement in expected returns alone from the NL weighting.

While our NL implementation is economically motivated, consistent with rational investor behavior, where past returns serve as a signal for expected returns ([Ferson and Siegel, 2001](#)), we compare our implementation versus empirically-driven weighting schemes. We examine linear signal weighting, the binary weighting scheme of [Moskowitz et al. \(2012\)](#) that simply goes long contracts with a positive past return and short those with a negative past return, and also an empirically-driven non-parametric weighting scheme using machine learning. Using machine learning techniques, we find that a neural net seeking to maximize the out-of-sample Sharpe ratio of the strategy uncovers a nonlinear weighting scheme that is very similar to the theoretical one we impose. The machine finds that the optimal signal weighting function is linear in the middle of the distribution and then flat or even moving slightly in the opposite direction in the extremes – just like the theoretical S-curve suggested by [Ferson and Siegel \(2001\)](#).²

²In addition to its gross outperformance of the linear and binary strategies, the NL strategy is also likely to improve after accounting for transactions costs, since the smooth nature of our nonlinear model reducing trading and extreme positions, in contrast to traditional binary or linear momentum weights. In binary momentum models in particular, a change in the sign of the forecast triggers a full reversal of the position from long to short (or vice versa), where small fluctuations in the forecast around zero can cause flip-flopping weights, resulting in frequent trading.

The rest of the paper is structured as follows. [Section 2](#) details our data and discusses the trend-following forecasts. [Section 3](#) describes our feasible, nonlinear time series momentum estimation. [Section 4](#) presents the empirical results and [Section 5](#) presents a number of robustness tests. [Section 6](#) concludes.

Literature Review. [Moskowitz et al. \(2012\)](#) explore time series momentum, where past 12-month returns of various asset classes – equities, commodities, currencies, and bonds – positively predict their next month’s returns, reversing over longer horizons. They suggest that initial underreaction and delayed overreaction may be driving these patterns, and that time series momentum provides substantial risk-adjusted returns, particularly during periods of high volatility. [Huang et al. \(2020\)](#) find that, while time series momentum is profitable, its predictability is not statistically robust when tested across a broader set of assets, especially when examining each asset class individually. [Goulding et al. \(2023\)](#) investigate how time series momentum strategies perform around market turning points. They differentiate between “slow” and “fast” momentum signals, using 12-months and 1-month lookback windows, respectively. Slow strategies provide better long-term results, but underperform around market reversals, when fast strategies are more effective. They suggest a blended “intermediate-speed” strategy combining both slow and fast momentum to better detect market turning points and balance responsiveness with stability. They identify four market cycles (Bull, Bear, Correction, Rebound), suggesting that this hybrid approach captures cycles more accurately. [Levine and Pedersen \(2016\)](#) apply a sign function to TSMOM and [Martin and Bana \(2012\)](#) illustrate the impact of various nonlinear mappings on the risk-adjusted performance and the skewness of returns of a TSMOM strategy.

Our paper extends this line of research by introducing nonlinear time series momentum motivated by simple investment theory. Our weighting scheme generates a flexible functional form (instead of a binary function), available at any frequency that is much more smooth and stable, resulting in lower trading and perhaps more feasible implementation.

2 Data

Our dataset comprises futures contracts from multiple asset classes, covering data through the end of October 2024 and beginning as early as January 1980. We consider front month futures for a set of 55 markets across different sectors. [Table 1](#) describes the contract specification. Most of the analyses in the paper use daily data, aggregated to weekly or monthly frequencies for some tests as well. Using daily returns offers the most granular and accurate representation of real-world dynamics for trend-following strategies.³

To generate a continuous series of prices for each market, all futures are generically rolled. The roll date is set as the earlier of the last trade date and the first day of the futures contract month. For equity futures we roll four days prior to each contract’s first notice date.⁴ This methodology provides a reasonable trade-off between liquidity and any unusual investor behavior that may occur in the last few days prior to the first notice date. For example, the May 2024 WTI Crude Oil futures contract could be traded until April 22, 2024, with first notice happening two days later. Thus, we roll this contract on April 22 into the corresponding June contract. Similarly, the June 2024 Mini-S&P 500 futures contract expired on June 21, 2024, so we roll it into the next expiration (September contract) four days earlier, on June 17, 2025.

On the dates when the rolling of the futures positions takes place, ratio adjustments are computed and applied sequentially backward in time. Starting from the most recent contract, the end-of-day prices of the previous contract are multiplied by the roll-adjustment factor, defined as

$$adj_t = \begin{cases} P_t^{far} / P_t^{near}, & \text{if } t \text{ is the roll date} \\ 1, & \text{otherwise} \end{cases} \quad (1)$$

where P_t^{near} is the price of the futures contract which will expire next and P_t^{far} is the price of the futures contract which will expire thereafter. As a clarifying example, for the Mini-S&P 500 futures roll on June 17, 2024, from the June 2024 contract to the September 2024

³A typical CTA engages in dynamic trading throughout the day, and CTAs actively engage in execution research (e.g., [Bordigonia et al., 2016](#)).

⁴The first notice date specifies the first time the clearing house may notify an investor that she has to make or take delivery of the underlying.

contract, the adjustment adj_t is defined as the last price of the September contract (P_t^{far}) on the 17th divided by the last price of the June contract (P_t^{near}) on the 17th. Using ratio adjustments for the rolls between contracts allows us to compute relative returns as is standard in the literature. Log-returns are then simply computed as the log-difference between the contiguous price. When we roll from the near contract (e.g., June) into the far one (e.g., July), the log return on that day is calculated as follows

$$r_t = \log(P_t^{far} / adj_t) - \log(P_{t-1}^{near}) = \log(P_t^{near}) - \log(P_{t-1}^{near}) \quad (2)$$

so that there is no price-jump from rolls but the carry (roll) return is earned.

In the return calculation (2) we assume that the futures position is fully collateralized, meaning that the return on the position is calculated on the notional amount. As an example, if a futures position of \$1mm notional amount with 20% initial margin (e.g., \$200,000) increases in value to \$1.2mm, the net return would be 20% (e.g., \$200,000/\$1,000,000), not 100%.⁵

Table 2 reports summary statistics of the futures in our sample. There is substantial heterogeneity in the returns to these contracts, many of which are used in trend-following and managed futures investment strategies.

3 Methodology

We begin with a primer on TSMOM signals and then introduce our NLTSOM signals that we analyze in the paper.

3.1 A Primer on Time Series Momentum

Time series momentum is an investment strategy that exploits positive autocorrelation in returns and hence can be described by a linear combination of past returns. A exhaustive in-depth review can be found in [Harvey et al. \(2021\)](#) and references therein. In a fairly

⁵We further assume there is no borrowing at the repo rate to fund the initial margin, and no repo interest on the initial margin received by the broker, e.g., the two interest rates offset each other.

general form, the TSMOM signal s_t can be defined as

$$s_t \equiv g(r_{t-\tau:t}) = k_1 \hat{\sigma}_t^{-1} \sum_{\tau=0}^{T-1} w_\tau r_{t-\tau} \quad \text{with} \quad k_1 \equiv \left(\sum_{\tau=0}^{T-1} w_\tau^2 \right)^{-1/2} \quad (3)$$

where g is any real-valued function, $r_{t-\tau}$ are lagged past returns of the asset, k_1 is a constant factor to normalize the signal to have unit variance, w_τ is a non-negative weight, and $\hat{\sigma}_t$ is a rolling one-year (e.g., 260 trading days) volatility estimate.⁶ This definition of TSMOM in (3) is simply that the signal s_t is a generic function of past returns.⁷

In practice, the lookback period T is typically chosen to be one year, as market return autocorrelation tends to reverse after that (Moskowitz et al., 2012). We choose T to be 1, 3, and 12 months in order to examine how the speed of the time series momentum signal affects the shape of the (nonlinear) forecast function.⁸ The weights in (3) can be set in various ways. For most of our analysis, we use a simple moving average, i.e., $w_\tau = 1/T$ (Hurst et al., 2017), in which case the normalizing constant becomes $k_1 = \sqrt{T}$. When g is linear, the data are monthly, and $T = 12$, we recover the standard linear trend-following momentum model:

$$s_t^{\text{Linear}} = r_{t-12m:t} / \hat{\sigma}_{t-12m:t}. \quad (\text{TSMOM-Linear})$$

which is widely used in the literature (Levine and Pedersen, 2016).⁹

To ensure that signals are comparable across asset classes and over time, two ap-

⁶We follow the common approach used in the literature and normalize, for example, the one-year return using a single volatility estimate. Alternatively, one can individually risk-adjust each past (e.g., daily) return, as in $\sum_{\tau=0}^{T-1} w_\tau r_{t-s} / \hat{\sigma}_{t-s}$. The (fixed) window size of 1 year is chosen for consistency. For example, using weekly data the volatility estimate for a one-month lookback would use only four observations if we matched the window size to the lookback.

⁷Abstracting from weights and focusing on the standard case where $T = 12$, a key problem with defining the signal $s_t = r_{t-1Y:t} = \sum_{\tau=0}^{11} r_{t-s}$ without scaling by volatility is that signals are not comparable across assets, e.g., they depend on the risk of the individual assets. For example, the signal has very different implications for a Treasury future compared to Natural Gas. This issue can be addressed by normalizing the signal by its volatility. This normalization has been applied in previous empirical tests of time series momentum (Moskowitz et al. (2012), Huang et al. (2020)).

⁸More in detail, we consider lookback horizons of 21, 62, and 260 trading days, for 1-month, 3-months, and 12-months, respectively. The analysis of different lookback periods is sometimes referred to as *momentum velocity*. Some time series momentum strategies combine different lookback periods (e.g., 1M-3M-12M), where shorter lookbacks increase the velocity of the signal.

⁹For daily data, (TSMOM-Linear) becomes $r_{t-260:t-1,t} / \hat{\sigma}_{t-260:t-1,t}$.

proaches have generally been adopted. The first approach is to scale the signal by its prevailing volatility estimate, as shown in equation (TSMOM-Linear). The second one is to use a binary signal function. A standard example in the literature is the binary trend-following signal introduced by Moskowitz et al. (2012),

$$s_t^{Binary} = \text{sign}(r_{t-\tau:t}) \quad (\text{TSMOM-binary})$$

where τ is the past τ months or trading days depending on the data frequency. The binary trend-following signal (TSMOM-binary) can be viewed as a special case of the linear trend-following signal in (TSMOM-Linear) where the volatility estimate is instantaneous $\hat{\sigma}_t = |r_t|$. This signal ignores the *magnitude* of the underlying movement.

In the following sections, we detail various momentum signals used in the literature and compare their performance to our NLTSMOM strategy. However, we first motivate theoretically why a nonlinear weighting of TSMOM is optimal, and then describe how we implement NLTSMOM in our tests.

In general, a nonlinear time series momentum trading strategy can be expressed as a function f of the signal s_t , e.g., $f(s_t) \equiv f[(g(r_{t-\tau:t}))]$. The functional form of the *optimal* function f can be imposed by theory or estimated from historical data given an optimality criterion.

3.2 A Simple Theoretical Framework for Nonlinear Time Series Momentum

From the perspective of an investor with mean-variance preferences, the optimal weighting of a signal will be a nonlinear function. For example, Ferson and Siegel (2001) consider a model with a riskless asset (with rate of return r_f) and a risky asset with return

$$R_t = \mu(s_t) + \varepsilon_t$$

where $\mu(s_t) = E(R_{t+1}|s_t)$ (where we denote R_{t+1} as the risk-adjusted return $R_{t+1} = r_{t+1}/\hat{\sigma}_{t+1|t}$). They show that the unique portfolio with minimum unconditional variance

for a given unconditional expected return places the following weight on the risky asset:

$$f_{FS}(s_t) = \lambda \frac{\mu(s_t)}{\mu(s_t)^2 + \sigma_\varepsilon^2}. \quad (4)$$

To implement the theoretical portfolio position $f_{FS}(s_t)$, we use the linear momentum signal defined in (3) and assume it perfectly predicts the next period's (risk-adjusted) return, implying $\mu(s_t) = s_t$. Given volatility scaling, a risk-normalized return is obtained by setting $\sigma_\varepsilon^2 = 1$, which simplifies the portfolio weight to

$$f_{FS}(s_t) = \frac{s_t}{s_t^2 + 1}.$$

The parameter λ in (4) is a constant which controls the risk aversion of the investor. To make the different signals across all strategies we examine comparable, we set λ such that $\text{var}(f_{FS}(s_t)) \approx 1$ for a standard normally distributed random variable s_t . This yields $\lambda \approx 0.394$.

Figure 1 plots the weighting function from [Ferson and Siegel \(2001\)](#) for a perfect forecast $\mu(s_t) = s_t$ and unit variance. As shown in the figure, as the magnitude of the signal, s_t , increases (decreases), the weight used on the signal increases (decreases) linearly between about -1 to 1 . Outside of that region, the signal weight starts decreasing (increasing) with the signal so that the weighting function is concave for $s_t > 1$ and convex for $s_t < -1$, where the signal has been standard normalized.

The intuition behind this nonlinear weighting is that when maximizing the unconditional Sharpe ratio is the objective, it is optimal to downweight extreme signals in order to make the tradeoff between maximizing returns and minimizing variance. Even when the signals contain no noise, extreme signals will predict returns more strongly, which will increase the weight on this signal, but the increased weight also makes the strategy more concentrated, resulting in less diversification and higher risk. At extreme signals, the increased risk effect dominates the increased return predictability and hence you get concave weights for large positive signals and convex weights for large negative ones.¹⁰

¹⁰If signals also contain noise, then the same logic would apply even if the objective were to maximize expected returns, because extreme signals are more likely to contain more noise and hence should be muted when forming optimal portfolio weights for out-of-sample predictability. A Bayesian framework would

In the empirical section we will see if the ranges of TSMOM signals reaches these non-linear regions and what impact that has on the out-of-sample performance of TSMOM strategies.

3.3 Empirical Nonlinear Time Series Momentum

While the [Ferson and Siegel \(2001\)](#) framework provides a simple nonlinear weighting scheme based on its assumptions, we can also explore other weighting schemes through machine learning that allow us to empirically find a weighting that meets some objective criterion. A common criterion is to minimize the mean squared forecast error (MSFE), $\sum_t (r_t / \hat{\sigma}_{t|t-1} - f(s_{t-1}))^2$ of the conditional return. We prefer an alternative criterion, which is to maximize the Sharpe ratio of the signal, e.g., $\sum_t r_t^s / (\sum_t r_t^{s2})^{0.5}$, where the proxy strategy return $r_t^s = f(s_{t-1}) r_t / \hat{\sigma}_{t|t-1}$ is used. Analogously, other measures such as the Sortino ratio could also be used.

In the empirical analysis, we impose symmetry on the function f that maps the signal s_t to the dollar trading position, e.g. $f(-s_t) = -f(s_t)$. While this constraint is not necessary, and we show that our results are robust to relaxing this assumption, it is sensible if we want the signal to be unbiased. Hence, if an unskewed, zero-mean random variable is passed into the system, the output will have mean zero. The signal in (3) satisfies this definition of unbiasedness, and imposing symmetry on f helps maintain this property.

We also want $f(0) = 0$, so that the strategy is market beta neutral. If $f(0) \equiv c \neq 0$, we can always rewrite the estimated function as $f(s_t) = \tilde{f}(s_t) + c$, where $\tilde{f}(0) = 0$. The residual constant strategy signal c would earn $c r_{t+1}$ and thus would resemble a buy and hold strategy on top of a pure TSMOM strategy. To keep comparisons across strategies clean and easy to interpret, we ensure all TSMOM strategies are beta neutral.¹¹

Finally, we also ensure that the function preserves the variance of the input, i.e. $\text{var}(X) \approx \text{var}(f(X))$, so that the application of this function to the forecast does not change the risk

naturally do this, for example.

¹¹Moreover, investors are generally unwilling (or should be) to pay trend-following fees for what is effectively a buy-and-hold exposure. An investor would prefer to buy a 'pure' trend-following strategy and, if she also wanted a constant exposure to the underlying market, buy a cheap exchange traded fund (ETF), for instance.

of the strategy unconditionally, for ease of comparison across strategies.¹² This condition merely applies a constant factor, which otherwise would be built into the leverage mechanism.

Next, we outline details on the estimation procedure. We estimate the function f using artificial neural networks (ANN) models.¹³ Given the amount of data we have, an artificial neural network regression provides a reasonable trade-off between flexibility and parsimony. We estimate these models with Keras (Chollet, 2015), performing a hyperparameter grid search on the training sample. The search is across the following hyperparameter dimensions:

1. the number of hidden layers for the ANN which is either 1 or 2;
2. the number of nodes in each layer which can be 2, 4, 8 or 16;
3. initial learning rate for the stochastic gradient descent algorithm (SGD) which can be 10^{-6} , 10^{-5} , 10^{-4} or 10^{-3} , which is tempered after 25 epochs using a learning rate scheduler that exponentially reduces the learning rate by $\exp(-0.05)$;
4. the weight on the L_2 -regularization for the parameters in each node which can be 10^{-4} , 10^{-3} , 10^{-2} or 10^{-1} ;
5. the drop rate for the learning is set to either 25% or 50%.

For all nodes we constrain the bias to be zero and choose hyperbolic tangent as activation functions. The output layer uses a linear activation in order to forecast risk-adjusted returns, which are unbounded. For training, we use a variant of the stochastic gradient descent known as Adam (Kingma and Ba, 2017). We use 50 epochs and split the sample for each epoch into batches of 1,024 samples each.

¹²By means of the volatility normalization of the original signal, data could be pooled across markets to estimate f . We study pooling in Section 5.2.2.

¹³In Section 5.2.1, we re-estimate our NLTSMOM function non-parametrically using the Nadaraya (1964) and Watson (1964) estimator (N&W).

For the training we use data up to 1994.¹⁴ For training, and validation sets, we symmetrize the data to ensure that the estimated function is symmetric. Specifically, we multiply the signal s_t by -1 , apply the function f estimated as described above, and then multiply the resulting strategy return by -1 . We minimize the negative Sharpe ratio as the objective function. Log returns are used to facilitate aggregation to lower frequencies via summation.

In the next subsection, we analyze the performance of a portfolio that trades nonlinear time series momentum using the above specified function f across a set of assets using futures contracts. We pool markets into three groups – equity indices, rates and FX, and commodities – and estimate the nonlinear function separately for each group.¹⁵ We then compare the NLTSMOM strategies to other TSMOM strategies from the literature.

3.4 Portfolio Returns

To deploy the same amount of risk for a given momentum forecast, positions are typically scaled inversely to the volatility of the underlying futures market. The dollar position $\text{pos}_{i,t}$ (for a \$1 portfolio) is determined by scaling the signal inversely to a forecast of volatility, and adjusting it to an arbitrary fixed target which we set to $\sigma_{\text{target}} = 12\%$ (this is a leverage factor to make the portfolios comparable), i.e.

$$\text{pos}_{i,t} = \frac{\sigma_{\text{target}}}{\hat{\sigma}_{i,t+1|t}} \cdot s_{i,t} \quad (5)$$

where i denotes the specific asset. This formulation also shows that the volatility scaling in the position formation is distinct from the signal normalization and that both are needed to have a risk-balanced portfolio of comparable trend-following signals.

The position (5) earns the returns $r_{i,t+1}$. In the linear case ([TSMOM-Linear](#)), assuming

¹⁴Specifically, the training ends in June 15, 1994 for daily, June 12, 1994 for weekly and June 30, 1994 for monthly data. We also experimented with an alternative training/validation split, where only every other observation is used in the training and every other observation is used for validation, and obtain similar (and sometimes better) results.

¹⁵Commodities include several sub-sectors: energies, agriculturals, precious metals, and base metals. We also tested a specification where we estimated only in each group, obtaining similar results.

a lookback period of $T = 1$ year, substituting $s_{i,t}^{\text{Linear}}$ into (5) gives:

$$\text{ret}_{i,t+1}^{TS-Mom} = \text{pos}_{i,t} r_{i,t+1} \propto \frac{r_{i,t-1Y,t}}{\hat{\sigma}_{i,t}^{1Y}} \frac{r_{i,t+1}}{\hat{\sigma}_{i,t+1|t}}. \quad (6)$$

The return earned by that signal is approximately given by today's signal applied to next period's return, scaled by today's forecast of next period's volatility.

When a nonlinear relationship exists, the dollar profit or loss (PnL) generated by the signal is approximately given by:

$$\text{ret}_{i,t+1}^{TS-Mom} \propto f(s_{i,t}) \frac{r_{i,t+1}}{\hat{\sigma}_{i,t+1|t}}. \quad (7)$$

The performance $\text{ret}_{i,t}^{TS-Mom}$ for each market $i = 1, \dots, N$ is calculated as in (7), and then aggregated (for a given sector or overall across all asset classes) using equal weights:

$$\text{ret}_t^{TS-Mom} = (1/N) \sum_{i=1}^N \text{ret}_{i,t}^{TS-Mom}. \quad (8)$$

Since the positions are scaled inversely to an estimate of underlying volatility, this approach effectively corresponds to volatility risk parity.

The Sharpe ratio for a sector is computed similarly by equally weighting strategy returns, where each strategy return is weighted in proportion to $1/\hat{\sigma}_{t+1|t}$.

4 Empirical Results

We present our empirical results on the performance of our nonlinear time series momentum strategy and compare it to other TSMOM strategies. We test for the statistical significance of Sharpe ratio differences by using the stationary bootstrap of [Politis and Romano \(1994\)](#).¹⁶

¹⁶We use $n_b = 1000$ bootstrap samples for each test. For each series of backtested returns the same set of indices is used, the Sharpe ratio calculated, and then the difference computed. We experimented with various dependency parameters for the stationary bootstrap with no qualitative difference in our results.

For each of the TSMOM signals we also estimate the predictive regression

$$r_{i,t+1}/\hat{\sigma}_{i,t+1|t} = \alpha_i + \beta_i s_{i,t} + \varepsilon_{i,t+1} \quad (9)$$

and compute the out-of-sample R^2 (Campbell and Thomson, 2008) defined as

$$R_{i,OS}^2 = 1 - \frac{\sum_{t=1}^T (r_{i,t} - \hat{r}_{i,t})^2}{\sum_{t=1}^T (r_{i,t} - \bar{r}_{i,t})^2} \quad (10)$$

which compares the mean-square error (MSE) of the prediction $\hat{r}_{i,t}$ with the MSE of the trailing mean estimate $\bar{r}_{i,t}$. The OOS R^2 reported is the average value across markets.

We examine the performance of our NLTSMOM strategies against various alternatives from the literature. Namely, the linear TSMOM signal, the Ferson and Siegel (2001) (FS) nonlinear signal from their model, and the binary signal of Moskowitz et al. (2012). We also make these comparisons across different observation frequencies, lookback periods, and horizons.

4.1 Frequency, Forecast Horizon, and Holding Period

4.1.1 Daily frequency

Table 3 shows results that combines using daily data and lookback periods of 1, 3, and 12 months (in Panels A-C, respectively). The first column of each panel reports the annualized Sharpe ratio of each TSMOM strategy. The next three columns report the p -values of pairwise tests of Sharpe ratio differences between the different weighting schemes. The last column reports the out-of-sample (OOS) R^2 .

In Panel A, for a lookback period of one month, our empirical NLTSMOM produces the highest Sharpe ratio of 0.83. The Ferson and Siegel (2001) nonlinear and the binary model produce similar Sharpe ratios of 0.59 and 0.53, respectively. All of these models outperform the linear model. However, the NL models are statistically no different from the binary model with a one month lookback.¹⁷

For the 3-month lookback horizon (Panel B), the Sharpe ratios for the nonlinear (0.47***)

¹⁷Table A.1 shows that these results are not driven by a specific period but hold consistently across our OOS period. Appendix Table C.1 shows similar conclusions for weekly lookback windows.

and empirical nonlinear (0.55***) models are significantly higher than that of the linear model (0.33*) with p -values below 0.08. Similarly, for the 12-month lookback horizon (Panel C), the Sharpe ratios for the nonlinear (0.83***) and empirical nonlinear (0.84***) models exceed that of the linear model (0.70***). Here, the empirical NLTSMOM is statistically different than the other models, and again the difference in Sharpe ratios between the linear and nonlinear models is statistically significant, with p -values below 0.03.¹⁸ Finally, note also that the out-of-sample improvement in Sharpe ratios is substantial, despite the fact that the out-of-sample predictive R^2 of the nonlinear models is only marginally higher than that of the linear model—though consistently positive.¹⁹

An investor seeking to build the best possible portfolio under mean-variance preferences will aim to maximize the risk-adjusted returns of her portfolio. Given the low correlation between Risk Parity (RP) and Time-Series Momentum (TSMOM) strategies,²⁰ a natural approach would be to start with a Risk Parity portfolio and then add TSMOM to enhance diversification and improve the overall Sharpe ratio. Table 4 shows the mean optimal weight to nonlinear TSMOM when combined with RP, highlighting that nonlinear momentum consistently commands a higher allocation than linear momentum across all lookback periods. The optimal weight to nonlinear TSMOM exceeds 50% in all cases, regardless of the specific nonlinear implementation. By contrast, the linear momentum strategy has a noticeably lower optimal weight than the nonlinear versions, with the gap narrowing slightly for longer lookback signals. This pattern suggests that nonlinear momentum strategies are more effective in improving portfolio efficiency, particularly with faster moving signals.

Table 5 drills down into specific asset classes, providing further evidence of a nonlinear momentum effect across all sectors and lookback periods. The last column shows

¹⁸The difference in Sharpe ratios between the empirical nonlinear and linear models is 0.14 for the 12-month lookback period and 0.22 for the 3-month lookback period. Interestingly, the smaller difference at the 12M horizon is statistically more significant. This is due to the lower standard deviation of returns associated with the 12M signal. Because the 12M signal evolves more slowly, the resulting strategy returns tend to be smoother.

¹⁹This finding resonates with those in Kelly et al. (2024).

²⁰For example, the strategy return correlations between linear TSMOM and RP are -0.25, -0.21 and 0.04 for the 1M, 3M, and 12M lookback periods. For the nonlinear TSMOM strategies, the correlations with RP are generally lower. In particular, for the empirical nonlinear momentum strategy, the correlations are -0.01, -0.07, and 0.01, highlighting a weaker association with RP across all lookback periods.

the p -values for the test of equivalence between nonlinear and linear momentum Sharpe Ratios. In most cases, the null hypothesis that nonlinear and linear models are equivalent can be rejected at the 10% significance level. The only notable exceptions are equities with a 1-month lookback period and commodities with a 3-month lookback period, where the differences are not statistically significant, though the NLTSMOM delivers more positive performance.²¹ Notably, the empirical nonlinear model often performs comparably to the binary model and, in several instances, outperforms it. For equities, the nonlinear model generally delivers a higher Sharpe ratio than the binary model for both the 1-month (0.13 vs. -0.02) and 3-month (0.17 vs. 0.03) lookback periods. This suggests that nonlinear effects are particularly valuable in capturing short-term ($<12M$) momentum signals.

Table A.1 presents results from an out-of-sample (OOS) analysis, where the function is estimated using data up to 2000, and the OOS period is divided into three subperiods. The findings suggest that the superior performance of nonlinear models relative to linear models is not limited to any specific subperiod. Focusing on the 12-month (12M) lookback period, both the Ferson and Siegel (2001) nonlinear model and the empirical nonlinear model consistently deliver higher Sharpe ratios than the linear model across all subperiods. Moreover, their Sharpe ratios are statistically different from zero at conventional significance levels in each of the three subsamples: 2000–2007, 2007–2015, and 2015–2024. For the 1-month (1M) and 3-month (3M) lookback periods, the Sharpe ratios of the linear model are not significantly different from zero, whereas the nonlinear models yield statistically significant Sharpe ratios, highlighting their superior ability to capture short-term momentum. It is also informative to compare nonlinear momentum with the binary model. The nonlinear models tend to outperform the binary model for shorter lookback signals (1M and 3M) but exhibit similar performance for longer lookback periods (12M). Finally, the decline in performance of the empirical nonlinear model in the

²¹The p -value tests the null hypothesis $H_0: \text{Sharpe}(\text{linear}) \geq \text{Sharpe}(\text{emp. nonlinear})$. For equities with a 1-month lookback period, the Sharpe ratio of the empirical nonlinear model exceeds that of the linear model. Notably, the linear model yields a negative Sharpe ratio in this case. Despite the apparent performance gap, the p -value remains above the conventional 0.1 threshold. This is largely driven by a sharp divergence in returns during the onset of the COVID-19 crisis in February 2020, when the empirical nonlinear model incurred losses while the linear model generated gains. The extreme return dispersion during that period increases the variance of the Sharpe ratio difference, thereby reducing the statistical significance of the result.

most recent period (2015–2024) may be attributed to the lack of re-estimation over time. In contrast, re-estimation would not affect the other methods, as their functional forms are fixed and not updated over the sample.

4.1.2 Monthly frequency

[Table 6](#) reports our main empirical results using monthly data instead of daily. At this lower frequency, the performance improvement of nonlinear models relative to the linear model is more muted compared to daily returns. Interestingly, however, the improvement in out-of-sample R^2 is more substantial, suggesting that nonlinear models offer greater predictive accuracy even when the gains in risk-adjusted returns are less pronounced. Moreover, even at this lower frequency, the nonlinear models, particularly the empirical nonlinear model, consistently outperform the linear model. For the classic 12-month (12M) momentum strategy, both the F&S and the empirical nonlinear models achieve the highest Sharpe ratio (0.65), outperforming the linear model as well as the alternative binary nonlinear specification. Statistical tests reject the null hypothesis of equal Sharpe ratios between the linear and F&S models (p -value = 0.09) and between the linear and empirical nonlinear models (p -value = 0.06), reinforcing the conclusion that nonlinear approaches more effectively capture long-term (12M) momentum signals—even at lower observation frequencies.

The improvement observed in [Table 6](#) for the classic 12-month (12M) lookback window is not driven by a specific period but holds consistently over time. [Table B.1](#) shows the results split across subperiods. Both in the 2000–2007 and 2007–2015 periods, the Sharpe ratio (SR) of the empirical nonlinear momentum model is significantly different from zero and larger than that of the linear model (1.16 vs. 0.87 in 2000–2007 and 0.69 vs. 0.54 in 2007–2015). This suggests that the nonlinear momentum effect is robust. However, in the most recent period (2015–2024), the Sharpe ratio for both the linear and nonlinear models is lower and not statistically different from zero. This pattern likely reflects the more challenging market environment for momentum strategies in recent years.

[Table 7](#) confirms and extends these findings by unpacking the results across different asset classes. It shows that while the improvement of nonlinear models over linear

models remains more subdued at a monthly frequency, there are still notable cases where nonlinear models deliver significant outperformance. For example, at the 3-month (3M) lookback period, the empirical nonlinear model generates a Sharpe ratio of 0.36 for equities, which is significantly higher than the linear model's Sharpe ratio of 0.25, with a p -value of 0.06—indicating that the improvement is statistically meaningful for equities. Moreover, for the classic 12-month (12M) momentum strategy at a monthly frequency, the empirical nonlinear model consistently delivers the highest Sharpe ratio among both the linear model and the nonlinear alternatives for equity and commodities. The gap between the empirical nonlinear and F&S models is relatively small, whereas it is larger between the empirical nonlinear and binary models. This suggests that, beyond nonlinearity, accurately capturing rollback effects for large signals may be important. These results underscore the strength of the empirical nonlinear approach in capturing longer-term momentum signals, even at lower observation frequencies.

[Table C.1](#) reports results at the weekly frequency. The table shows that both the F&S nonlinear and empirical nonlinear models deliver performance comparable to the binary model and statistically superior to the linear model for both short (1M) and long (12M) lookback periods. The differences between the F&S nonlinear and linear models, as well as between the empirical nonlinear and linear models, are statistically significant, with p -values of 0.02 and 0.06, respectively.

[Table C.2](#) shows that the empirical nonlinear model is the only one to deliver positive Sharpe ratios for equities at the 1-month (1M) lookback period. Interestingly, for equities, the binary nonlinear model generates negative Sharpe ratios at both the 1M and 3M horizons. The benefits of nonlinear momentum are particularly strong for rates and FX. Specifically, we reject the null hypothesis of equal Sharpe ratios between the empirical nonlinear and linear models at both the short (1M) and long (12M) lookback periods. For commodities, the impact of nonlinearities becomes more pronounced at longer horizons.

[Figure 2](#) displays the estimated nonlinear function for the three different lookback periods (1M, 3M, and 12M) across various observation frequencies (daily, weekly, and monthly) for equity futures. The results suggest that the shape of the nonlinear function is more sensitive to the observation frequency than to the lookback period. The em-

pirically estimated function appears fairly linear for momentum signals in the range of -1 to 1 (around the origin), consistent with [Ferson and Siegel \(2001\)](#)’s theoretical result. However, for larger momentum signals, the function exhibits a rollback effect, which resembles a form of robust control or shrinkage. At higher frequencies (e.g., weekly), the rollback effect is more pronounced for shorter lookback periods (1M) than for longer ones (12M), suggesting that short-term signals are more sensitive to extreme momentum values when measured at high frequency. In contrast, at lower frequency (monthly), the pattern reverses: the rollback is more pronounced for longer lookback periods (12M) compared to shorter ones (1M and 3M).

Overall, the nonlinear models significantly outperform the linear TSMOM models across a variety of horizons, lookback periods, and data frequencies, as well as across and within the asset classes. We next investigate where this outperformance might be coming from by investigating another key feature of TSMOM strategies – their tail hedging properties.

4.2 Tail Hedging and Convexity

One of the main features of TSMOM strategies is their convexity and tail hedging properties with respect to the equity market portfolio. [Moskowitz et al. \(2012\)](#) were the first to show that TSMOM returns are convex with respect to equity market returns – performing better on average when the stock market (CRSP VW index or MSCI World index) performs extremely well or extremely poorly. [Hurst et al. \(2017\)](#) replicate this feature of TSMOM out of sample over a century and show that TSMOM does especially well when the stock and bond markets crash – providing tail hedging to stock and bond markets. Indeed, one of the major selling points of trend following managed futures strategies to investors is that they offer this tail hedge while simultaneously producing abnormal positive alpha unconditionally.

In this subsection, we investigate how our NLTS MOM strategies do in these market environments and address whether their additional unconditional performance over linear TSMOM maintains this hedging property, diminishes it, or improves it, and, if the latter, whether this tail performance is primarily driving the overall outperformance of

NLTSMOM. This assessment is also a test of the underlying theory for why a nonlinear weight might matter. Since the idea is to conditionally mute more extreme signals in order to maximize the unconditional Sharpe ratio, this muting effect would seem to be most important during the most extreme market environments. In particular, during extreme market downturns, when volatility typically spikes and financing may be expensive, a more robust weighting scheme designed to dampen extreme weights may be particularly valuable.

[Table 8](#) reports results at the monthly frequency for a 1-month forecasting horizon, using 12-month momentum as the signal across different market environments. We classify market regimes based on monthly returns on the Mini S&P 500 futures contract. Specifically, we define high market states as months in which the return exceeds 1.33%, low states as months in which the return falls below -1.33% , and medium states as those with absolute returns less than or equal to 1.33% . This threshold corresponds to an annualized average return of 16% and results in approximately equal numbers of observations across the three buckets. The unconditional regime includes all observations and serves as the baseline comparison. As a robustness check, we also employ tighter definitions of market states using a 20–60–20 split, where the bottom and top states each contain 20% of the observations.

[Table 8](#) shows that the nonlinear momentum strategy with a 12-month lookback and a 1-month forecasting horizon outperforms the linear strategy not only unconditionally, but also within specific market regimes. In particular, the performance gain is concentrated in bad market states—with p -values of 0.00 across all nonlinear specifications—and, to a lesser extent, in medium states.²² The advantage of the nonlinear approach is largely state-dependent. [Table D.1](#) confirms these conclusions when market regimes are defined more tightly using a 20–60–20 split.

The previous analysis demonstrates that nonlinear strategies with a 12-month lookback outperform their linear counterparts unconditionally and in low market states when using a 1-month forecasting horizon. [Table 9](#) presents results for a 3-month holding pe-

²²In the medium state, only our nonlinear empirical strategy outperforms the linear benchmark at the 10% significance level. By contrast, the FS and binary implementations do not yield significant gains. This conclusion is robust to the choice of threshold used to define market regimes; see [Table D.1](#).

riod, and [Table 10](#) reports results for a 12-month holding period. For each alternative horizon, we re-estimate the nonlinearity.²³ Despite the longer holding periods, signals continue to be updated monthly.

[Table 9](#) shows that the nonlinear momentum strategy continues to outperform the linear strategy, particularly in bad and medium market states, with p -values of 0.02 and 0.06, respectively, for our empirical nonlinear specification. Interestingly, the binary strategy delivers a statistically significant improvement over the linear strategy only in bad market states (p -value = 0.00), but not in medium states (p -value = 0.52). This suggests that the nonlinear specification captures additional predictive content beyond what is captured by a simple binary filter, especially in more moderate market environments. [Table D.2](#) confirms these conclusions when using the tighter 20–60–20 regime split.²⁴

The previously documented advantages of nonlinear momentum strategies dissipate at the 12-month forecasting horizon, as shown in [Table 10](#). At this longer horizon, the linear and nonlinear strategies perform similarly, both unconditionally and across market states.²⁵

The outperformance of our nonlinear strategy relative to the linear benchmark for 1-month and 3-month holding periods is also evident in the graphical analysis. [Figure 3](#), Panel A, displays the estimated nonlinear functions for equities across 1-month, 3-month, and 12-month forecasting horizons. As shown, nonlinearities diminish as the forecasting horizon increases, consistent with our earlier findings that the Sharpe ratio differences are statistically significant for shorter horizons (1M and 3M), but not for longer ones. In contrast, Panel B of [Figure 3](#) presents the corresponding estimates for FX and rates. In these asset classes, nonlinearities appear to become more pronounced with longer holding pe-

²³As in the baseline analysis, returns over the holding period are scaled using the ex-ante volatility estimate available at the time of signal formation. For example, a 12-month holding period from January to December 2023 uses the volatility estimate from December 2022. All strategies are scaled to target an annualized volatility of 12%.

²⁴[Table 9](#) and [Table D.2](#) show that the outperformance of our nonlinear empirical strategy relative to the linear benchmark is statistically significant at the 10% level in the medium state. By contrast, the FS and binary implementations do not exhibit significant gains.

²⁵We also conducted analysis using longer lookback windows of 36 months. For the 36M we exclude the first year, i.e. the signal in Dec 2022 is based on the return from January 2020 to December 2021 (both inclusive). However, we did not find any statistically significant effects for any TSMOM strategy and no significant differences for nonlinear versus linear strategies at this longer lookback, across all holding periods considered (1-month, 3-month, and 12-month). We omit these results from the tables.

riods. This divergence suggests that a sector-by-sector analysis may yield additional insights and further improvement in performance.

Finally, [Table 11](#) shows that our nonlinear momentum strategy generates statistically significant CAPM alphas across various holding periods (see separate panels) and look-back windows (see rows within each panel), both unconditionally and conditional on market states (see columns within each panel). The only exception is the 36-month look-back, which generally fails to produce statistically significant alphas of any kind for any TSMOM method.

Overall, a significant part of the additional performance attributed to NLTSMOM is its performance during extreme market environments, particularly equity market downturns. This finding is consistent with the theoretical result of a nonlinear weighting scheme maximizing the Sharpe ratio of an investment strategy. Both the theoretical nonlinear weights of [Ferson and Siegel \(2001\)](#) and our empirical nonlinear weights from machine learning, deliver better unconditional performance during these tail events than the linear TSMOM specification or the binary specification. The bulk of the outperformance of the NLTSMOM strategies seems to come from the conditional returns in bad market states. However, our empirical NLTSMOM strategy also outperforms in moderate states, suggesting that our machine learning approach also captures additional insight that improves TSMOM predictability even outside of bad market states. NLTSMOM delivers better tail hedging properties and better overall performance than traditional TSMOM strategies, consistent with theory.

5 Robustness Tests

In this section, we examine how robust our NLTSMOM model is by applying it to a unique set of assets – equity long-short factors – as an OOS test, using different estimation methods, and also comparing our results to XSMOM strategies among the same set of assets.

5.1 Equity Factor Momentum

[Gupta and Kelly \(2019\)](#) and [Ehsani and Linnainmaa \(2022\)](#) document factor momentum – positive autocorrelation in long-short equity factors from the academic literature – that exhibit trend following behavior. We examine the robustness of our nonlinear weighting scheme by applying it to these equity factors and compare our method to the simple linear trends used in the literature.

[Table 12](#) and [Table 13](#) present the results of our analysis employing the Fama–French factors SMB, HML, RMW, and CMA, for one-day and one-month holding periods, respectively, across 1, 3, and 12-month lookback periods.

[Table 13](#) shows that at the 3-month and 12-month lookback horizons, both the FS and empirical nonlinear strategies achieve Sharpe ratios that are statistically higher than the linear benchmark at conventional levels. By contrast, the binary strategy does not exhibit significant improvements over the linear specification. Interestingly, when we turn to the daily frequency in [Table 12](#), all nonlinear methods, including the binary, deliver Sharpe ratios significantly higher than the linear benchmark for short lookback periods of 1 and 3 months, but the differences are not statistically significant at the longer 12-month horizon.

These results provide another out-of-sample test of the nonlinear model to a set of completely different assets – long-short market neutral equity factors – that showcase the robustness of our NLTSMOM model.

Having shown that our nonlinear model performs well across the standard Fama–French factors, we next expand the analysis to a substantially larger universe by examining the 153 factors from [Jensen et al. \(2023\)](#). In this broader setting, we concentrate on the more common one-month holding period and assess whether similar nonlinearities arise. In particular, we cluster the 153 factors from [Jensen et al. \(2023\)](#) into the 13 themes defined in their study. For the theme-level statistics, we aggregate all (ex-ante) volatility-scaled factors within each theme using equal weights, consistent with our treatment of the futures universe in the earlier analysis. We also report results for an aggregate portfolio that combines all themes, again using equal-weighted, volatility-scaled factors. To ensure a complete cross-section for in-sample estimation, we begin the sample in 1980 and use

the period starting in January 2000 for out-of-sample evaluation.

Figure 4 reports the nonlinearities across these equity factors for 1-, 3-, and 12-month lookback horizons under a one-month holding period. Figure 5 compares the Sharpe ratios of the linear and empirical nonlinear specifications across the 13 factor themes for the same set of lookback windows (one per row). Each bar within a row corresponds to a theme, while the two rightmost bars display the linear and nonlinear Sharpe ratios for the aggregate portfolio that pools all factors irrespective of theme. Statistical significance is assessed using the bootstrapped difference between the empirical nonlinear and linear specifications and is reported above the empirical nonlinear estimates.

For the three-month lookback window, nonlinear momentum applied at the factor level generally delivers higher Sharpe ratios than its linear counterpart with statistically significant improvements in 7 out of the 13 themes. The differences are statistically significant for several widely studied anomalies, including size, value, investment, and, to a lesser extent, profit growth.²⁶ Interestingly, some categories such as profitability and seasonality, which do not show significant differences at the three-month horizon, do exhibit statistically significant Sharpe improvements at the 1-month lookback period.

When pooling all 153 anomalies using inverse-volatility weights (the last two bars in each plot), the nonlinear specification yields statistically significant Sharpe ratios for both the one-month and three-month windows. By contrast, results based on the twelve-month lookback are generally insignificant, suggesting that the performance advantage of nonlinear momentum for equity factors seems concentrated at shorter-horizons.

5.2 Specification Robustness

We assess the robustness of changing some of the specifications of our estimated model. Specifically, we examine the impact of removing data symmetrization and the impact of pooling data from all markets rather than only from sub-groups.

²⁶These findings align with the patterns observed for the Fama–French factors (see Table 13), where the empirical nonlinear specification at the monthly frequency performs best at the three-month lookback.

5.2.1 Non-parametric Estimation

As a robustness test, we estimate the function non-parametrically using the [Nadaraya \(1964\)](#) and [Watson \(1964\)](#) estimator (NW). The results for equities are shown in [Figure 6](#), estimated over the same in-sample period as [Figure 2](#), (e.g., up to 1999). We use an Epanechnikov kernel and the bandwidth is set to $b = 1$ as the input data is volatility-normalized.²⁷ Of note is that the objective function here is different to the one employed in the previous sections, which explicitly maximize the expected Sharpe ratio. As the figure shows, the results are largely consistent with our previous findings.

5.2.2 Pooling

Previously, we estimated the nonlinear function within sectors (Equities, Rates and FX, and Commodities). However, with the volatility normalization of the original signal and the regressand, data can be pooled across markets to jointly estimate the function, f . [Figure 7](#) reports the estimated nonlinear function when we pool the data from all markets to estimate a single forecast function.

We continue to observe clear evidence of nonlinearities and rollback effects for signals with large absolute values, suggesting that these features are not asset class-specific but instead represent a general property of time series momentum forecasts. Again, the results remain largely consistent under this alternative specification for estimating f .

5.3 XSMOM and Other Timing Strategies

Finally, we compare our NLTSMOM strategies to XSMOM strategies applied to the same assets. We then take some of the innovations and improvements the literature has found for XSMOM and apply them to TSMOM to see if our nonlinear weighting adds value over, or is picking up, these nuanced measures of momentum.

²⁷We found that simple cross-validation tends to select a very small bandwidth which ultimately leads to fairly unstable estimates.

5.3.1 Comparison to XSMOM

Table 14 compares Sharpe ratios of time-series and cross-sectional momentum strategies across different lookback periods. While cross-sectional momentum is typically implemented with a 12-month lookback, we also examine shorter horizons (1M and 3M), which have not been studied previously. For reference, the Sharpe ratios of the time-series strategies are the same as those reported in Table 6, with the addition here of the 36-month lookback. At this longer horizon, however, none of the strategies deliver statistically significant results.

First, Panel A shows that cross-sectional momentum delivers economically meaningful Sharpe ratios at the standard 12-month horizon (0.48), but it is smaller than the Sharpe ratios achieved by nonlinear time-series momentum – 0.65 Sharpe ratio for FS and empirical NL, and 0.61 for the binary strategy – the latter result also noted by Moskowitz et al. (2012). At shorter horizons, cross-sectional momentum is only marginally significant at 1M and not significant at 3M, which may justify the literature’s focus on longer lookback periods. However, TSMOM is quite significant at 1 and 3 month lookback periods.

To assess the incremental value of time-series momentum, we compute appraisal ratios of the various time-series strategies relative to cross-sectional momentum at the corresponding lookback period. The results, reported in Panel B, show that both linear and nonlinear time-series momentum deliver significant appraisal ratios relative to cross-sectional momentum at the 3-month and 12-month horizons. For instance, at the 12-month horizon, moving from cross-sectional to nonlinear (FS or empirical) time-series momentum raises the Sharpe ratio from 0.48 to 0.65. This gain is equivalent to combining the cross-sectional strategy with an orthogonal strategy of Sharpe ratio $\sqrt{0.65^2 - 0.48^2} = 0.44$, which coincides with the Treynor–Black appraisal ratio.

This performance improvement can be decomposed into two components: one due to volatility scaling and the other due to mean forecasting. If we consider only volatility scaling, the appraisal ratio is $\sqrt{0.55^2 - 0.48^2} = 0.27$, which is marginally significant at the 10% level. The remainder—captured by mean forecasting—provides a substantial incremental gain. Thus, while volatility scaling explains part of the improvement, an

important source of outperformance comes from better forecasting expected returns using nonlinear specifications.

Importantly, the significant appraisal ratios of time-series momentum strategies relative to cross-sectional momentum persist when the holding period is extended to three months, as shown in [Table 15](#). At the 12-month lookback, for example, moving from cross-sectional to nonlinear (FS or empirical) time-series momentum raises the Sharpe ratio from 0.73 to 1.03. This improvement is equivalent to combining the cross-sectional strategy with an orthogonal strategy that has a Sharpe ratio of $\sqrt{1.03^2 - 0.73^2} = 0.73$, which matches the Treynor–Black appraisal ratio. Moreover, the increase in appraisal ratio—from 0.44 at the 1-month holding period to 0.73 at the 3-month holding period—may reflect the weaker persistence of cross-sectional rankings compared to time-series signals.

5.3.2 Comparison with Other Timing Strategies

Finally, we examine other timing strategies inspired by improvements researchers have discovered with regard to XSMOM and apply them in a TSMOM context.

Starting with the linear trading signal,

$$s_t^{\text{linear}} = \frac{r_{t-12:t}}{\hat{\sigma}_{t-12:t}}, \quad r_t^{\text{TSMOM}} = s_{t-1}^{\text{linear}} r_t$$

We use the [Barroso and Santa-Clara \(2015\)](#) constant volatility (C-Vol) strategy to time the assets. [Barroso and Santa-Clara \(2015\)](#) find that constant volatility weighting has a marked improvement on cross-sectional equity momentum strategies, and [Daniel and Moskowitz \(2016\)](#) find that this weighting scheme also improves XSMOM in other asset classes (equity index futures, fixed income, currencies, and commodities) as well. To implement [Barroso and Santa-Clara \(2015\)](#) in a time-series context, we define

$$r_t^{\text{BSC}} = s_{t-1}^{\text{linear}} r_t \frac{\sigma_{\text{target}}}{\hat{\sigma}_{t|t-1}}.$$

We also examine a strategy that scales by variance,

$$r_t^{\text{var}} = s_{t-1}^{\text{linear}} r_t \frac{\sigma_{\text{target}}^2}{\hat{\sigma}_{t|t-1}^2}$$

as a modified version of [Barroso and Santa-Clara \(2015\)](#) applied to TSMOM.

[Daniel and Moskowitz \(2016\)](#) offer a dynamic weighting function to time XSMOM strategies within asset classes that weights a long-short XSMOM strategy by an ex ante estimate of its conditional Sharpe ratio, where they use a GARCH model to predict the conditional volatility of the strategy and a regression model based on the prior two years of returns on the market interacted with market volatility to predict the conditional mean of the XSMOM strategy. Applied to TSMOM, this takes the following form,

$$r_t^{\text{DM}} = s_{t-1}^{\text{linear}} r_t \frac{\hat{\mu}_{t-1}}{\hat{\sigma}_{t|t-1}^2}$$

where

$$\hat{\mu}_t = \hat{\gamma}_{0,t} + \hat{\gamma}_{1,t} \hat{\sigma}_t^2 \mathbf{1} \left(r_{t-24:t}^{\text{SPX}} < 0 \right),$$

which is estimated by OLS in the expanding window regression:

$$r_t^{\text{TSMOM}} = \gamma_{0,t} + \gamma_{1,t} \hat{\sigma}_t^2 \mathbf{1} \left(r_{t-24:t}^{\text{SPX}} < 0 \right) + \varepsilon_t. \quad (11)$$

We then compare each of these modified TSMOM strategies to the linear, binary, and our NLTSMOM strategies (theoretical and empirical) used throughout the paper.

Our NLTSMOM strategies essentially adopt the [Barroso and Santa-Clara \(2015\)](#) weighting scheme but replace the linear signal with a nonlinear signal as the base strategy.

In [Barroso and Santa-Clara \(2015\)](#), each asset is scaled by its volatility. To implement the DM timing strategy, we estimate the regression in (11) for each asset and compute the weight $\hat{\mu}/\hat{\sigma}^2$ at the asset level. The expanding regression used to construct the DM weights begins with two years of data, ensuring that the bear-state indicator is well defined, and then extends through the end of the sample in September 2024. In practice, we use the variance-scaled TSMOM returns as the basis and multiply them by the mean estimate $\hat{\mu}$ from the regression. All comparisons are restricted to the 12-month lookback and

1-month holding period, ensuring consistency with the existing literature. Finally, unlike Barroso and Santa-Clara (2015) and Daniel and Moskowitz (2016) we estimate risk using a 12-month window to match the return lookback periods. The prior studies instead rely on a mixed-frequency approach, calculating volatility or variance with a 6-month window.

Table 16 reports Sharpe ratios and pairwise appraisal ratios for alternative TSMOM strategies over the period 2000 to 2024. The Sharpe ratios for the linear (constant-volatility, following Barroso and Santa-Clara (2015)) and nonlinear strategies match those reported in Table 6–Panel C. The Sharpe ratios are generally higher for nonlinear strategies compared to variance scaling or the dynamic DM strategy. Focusing on the appraisal ratios, the fourth row compares the nonlinear strategies to the dynamic regression benchmark of Daniel and Moskowitz (2016). Moving from the dynamic regression to either the FS or the empirical nonlinear strategy raises the Sharpe ratio from 0.25 to 0.65. This improvement is equivalent to combining the dynamic strategy with an orthogonal strategy that has a Sharpe ratio of $\sqrt{0.65^2 - 0.25^2} \approx 0.61$, which corresponds to the Treynor–Black appraisal ratio. The performance gain can be decomposed into two components: one due to volatility scaling (rather than variance scaling) and the other due to nonlinear mean forecasting.

The second row compares the nonlinear strategies to the constant-volatility benchmark of Barroso and Santa-Clara (2015). Here, the binary strategy delivers the smallest appraisal ratio (0.28), significant only at the 10% level, whereas both the FS and empirical nonlinear strategies achieve appraisal ratios above 0.38 that are statistically significant at the 5% level.

Comparing the simple linear strategy to the FS or empirical nonlinear strategy shows a substantial improvement: the Sharpe ratio rises from 0.37 to 0.65, equivalent to adding an orthogonal strategy with a Sharpe ratio of $\sqrt{0.65^2 - 0.37^2} \approx 0.53$. This increase can be decomposed into two parts: the effect of volatility scaling (appraisal ratio = 0.45) and the incremental gains from nonlinear mean forecasting (appraisal ratio = 0.38 for FS and 0.41 for the empirical nonlinear strategy, relative to the constant-volatility benchmark). Since both the constant-volatility and nonlinear strategies are already volatility-scaled, the additional improvement from 0.55 to 0.65 directly reflects the contribution of nonlin-

ear mean forecasting, which is substantial.

Overall, the NLTSMOM strategies, motivated by theory, consistently outperform the linear, volatility and variance scaled, and dynamically weighted TSMOM strategies. The fact that the machine learning algorithm is able to construct a nonlinear weighting scheme that matches theory and delivers superior performance over linear and other forms of TSMOM, both unconditionally and during market downturns, is interesting and consistent with previous results. These findings being unique and distinct from XSMOM is also interesting and further emphasizes that TS and XS momentum are related, but different phenomena. Whether our nonlinear weighting could improve upon XSMOM remains to be seen and is left for future research.

6 Conclusion

We document robust empirical evidence of a nonlinear relationship between past returns signals and next-period expected risk-adjusted returns that is consistent with theory. Nonlinear time series momentum consistently improves out-of-sample performance across asset classes and time periods. The nonlinear strategies not only outperform the linear strategy on average over the full sample period, but also maintain, and improve upon, the convexity or tail-hedging property of TSMOM, where returns are highest during market downturns. NLTSMOM's outperformance is largely driven by the added returns during these downturns.

These results are remarkably stable across a range of estimation approaches (non-parametric or parametric, holds across return frequencies (daily, weekly, and monthly), with stronger effects at higher frequencies, for different horizons and lookback periods, and across different test assets, including long-short equity factors. Importantly, the shape of the nonlinear function is preserved regardless of the neural network's loss function, confirming that the result is not an artifact of a specific optimization criterion. The non-linearity typically takes the form of a rollback function, dampening extreme signals that match theoretical results.

These results carry two important implications: First, a simple nonlinear transforma-

tion of a momentum signal can lead to statistically significant improvements in strategy performance, both unconditionally and conditional on market downturns. Second, the estimated nonlinear function offers deeper insight into the behavior and limitations of conventional linear momentum strategies, highlighting how signal strength interacts with predictability in a nonlinear fashion that matches theoretical predictions.

References

- Asness, Clifford S. (1994) "Variables that explain stock returns," *Ph.D. Dissertation. University of Chicago*.
- Babu, A., A. Levine, Y.H. Ooi, L.H. Pedersen, and E. Stamelos (2020) "Trends Everywhere," *Journal of Investment Management*, 18 (1), 52–68.
- Baltas, N. (2019) "The Impact of Crowding on Alternative Risk Premia Investing," *Financial Analysts Journal*, 75 (3).
- Barberis, Nicholas, Andrei Shleifer, and Robert Vishny (1998) "A model of investor sentiment," *Journal of Financial Economics*, 49 (3), 307–343.
- Barroso, P. and P. Santa-Clara (2015) "Momentum and Its Moments," *Journal of Financial Economics*, 116, 111–120.
- Bollen, N.P.B., M. Hutchinson, and J. O'Brien (2021) "When it pays to follow the crowd: Strategy conformity and CTA performance," *Journal of Futures Markets*, 41 (6).
- Bordigonia, Giuliana, Alessio Figalli, Anthony Ledford, and Philipp Ustinov (2016) "Strategic Execution Trajectories," *Applied Mathematical Finance*, 29 (4), 288–330.
- Boudoukh, Jacob, Yukun Liu, Tobias J. Moskowitz, and Matthew P. Richardson (2025) "Identifying Shocks to Systematic Risk in Times of Crisis," *Working Paper*.
- Campbell, J.Y. and S.B. Thomson (2008) "Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?" *The Review of Financial Studies*, 21, 1509–1531.
- Chollet, F. (2015) "Keras," Technical report, <https://github.com/fchollet/keras>.
- Daniel, K. and Tobias J. Moskowitz (2016) "Momentum Crashes," *Journal of Financial Economics*, 122.

- Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam (1998) "Investor Psychology and Security Market Under- and Overreactions," *The Journal of Finance*, 53 (6), 1839–1885.
- Ehsani, Sina and Juhani T. Linnainmaa (2022) "Factor Momentum and the Momentum Factor," *The Journal of Finance*, 77 (3), 1877–1919.
- Person, W.E. and A.F Siegel (2001) "The Efficient Use of Conditioning Information in Portfolios," *Journal of Finance*, 56 (3), 19–38.
- Forner, C., Y.M. Muradoglu, and S. Sivaprasad (2018) "Enhancing Momentum Investment Strategy using Leverage," *Journal of Forecasting*, 37 (5).
- Goulding, Christian L., Campbell R. Harvey, and Michele G. Mazzoleni (2023) "Momentum turning points," *Journal of Financial Economics*, 149 (3), 378–406.
- Gupta, Tarun and Bryan Kelly (2019) "Factor Momentum Everywhere," *The Journal of Portfolio Management*, 45 (3), 13–36.
- Harvey, C.R., S. Rattray, and O. van Hemert (2021) *Strategic Risk Management: Designing Portfolios and Managing Risk*, Hoboken, New Jersey: Wiley.
- Hong, Harrison and Jeremy C. Stein (1999) "A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets," *The Journal of Finance*, 54 (6), 2143–2184.
- Huang, Dashan, Jiangyuan Li, Liyao Wang, and Guofu Zhou (2020) "Time series momentum: Is it there?" *Journal of Financial Economics*, 135 (3), 774–794.
- Hurst, B., Y.H. Ooi, and L.H. Pedersen (2017) "A Century of Evidence on Trend-Following Investing," *Journal of Portfolio Management*, 44 (1).
- Jegadeesh, Narasimhan and Sheridan Titman (1993) "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency," *The Journal of Finance*, 48 (1), 65–91.

- Jensen, Theis Ingerslev, Bryan Kelly, and Lasse Heje Pedersen (2023) "Is There a Replication Crisis in Finance?" *The Journal of Finance*, 78 (5), 2465–2518, <https://doi.org/10.1111/jofi.13249>.
- Kelly, Bryan, Semyon Malamud, and Kangying Zhou (2024) "The Virtue of Complexity in Return Prediction," *Journal of Finance*, 79 (1), 459–503.
- Kelly, Bryan T., Tobias J. Moskowitz, and Seth Pruitt (2021) "Understanding momentum and reversal," *Journal of Financial Economics*, 140 (3), 726–743.
- Kim, A.Y. and J.K. Tse, Y. Wald (2016) "Time Series Momentum and Volatility Scaling," *Journal of Financial Markets*, 30.
- Kingma, D.P. and J. Ba (2017) "Adam: A Method for Stochastic Optimization," Technical report, arXiv: 1412.6980.
- Kwon, O.K. and S. Satchell (2021) "Treating cross sectional and time series momentum returns as forecasts," *Journal of Forecasting*, x (x).
- Levine, Ari and Lasse Heje Pedersen (2016) "Which Trend Is Your Friend?" *Financial Analysts Journal*, 72 (3), 51–66.
- Martin, Richard and Ali Bana (2012) "Non-linear Momentum Strategies," *Risk*, 60–65.
- Moskowitz, Tobias J., Yao Hua Ooi, and Lasse H. Pedersen (2012) "Time Series Momentum," *Journal of Financial Economics*, 104 (2), 228–250.
- Nadaraya, E.A. (1964) "On Estimating Regression," *Theory of Probability and Its Applications*, 9 (1).
- Politis, D.N. and J.P. Romano (1994) "The Stationary Bootstrap," *Journal of the American Statistical Association*, 89 (428).
- Watson, G.S. (1964) "Smooth regression analysis," *Sankhya: The Indian Journal of Statistics, Series A*, 26 (4).

Zakamulin, V. and J Giner (2020) "Trend Following with Momentum Versus Moving Average: A Tale of Differences," *Quantitative Finance*, 20 (6).




Figure 1: Ferson and Siegel (2001) nonlinearity. This figure shows the nonlinearity for perfect forecasts, i.e. $\mu(s_t) = s_t$ and a unit variance error term, under the [Ferson and Siegel \(2001\)](#) framework.




Figure 2: Nonlinear Time Series Momentum - equities. This figure plots the linear TSMOM signal on the x-axis and the corresponding trading forecast on the y-axis for all equity futures jointly. The functions is estimated using artificial neural networks. The sample period is from January 1980 (where available) to December 1999.




Figure 3: Forecast function as a function of signal strength. This figure presents the estimated forecast functions based on a 12-month lookback window for three sectors: equity futures (Panel A), rates and FX (Panel B), and commodities (Panel C). Within each panel, we report the estimated functions for three alternative holding periods. The sample period spans from January 1980 (where available) to December 1999.




Figure 4: Non linearities in equity factors. This figure displays the nonlinearities in the equity factors of [Jensen et al. \(2023\)](#), clustered into the 13 themes defined in their study. The in-sample period begins in 1980 to ensure a balanced cross-section, and the out-of-sample period spans January 1, 2000 to December 31, 2024.




Figure 5: Sharpe ratio comparison: equity factors. This figure shows the Sharpe ratios for the linear TSMOM vs. the empirical nonlinear TSMOM. The top panel reports the comparison for the 1-month lookback period, the middle panel for the 3-month lookback period, and the bottom panel for the 12-month lookback period. Each theme is equally risk weighted (with ex-ante vol). Statistical significance is based on the bootstrapped difference between empirical nonlinear and linear. OOS period: January 1st, 2000 to December 31, 2024.




Figure 6: Forecast function as a function of signal strength. This figure shows the estimated forecast function using a traditional local constant kernel-density estimator for equities for 1M, 3M and 12M lookback for different data frequencies for equity futures. The sample period is from January 1980 (where available) to December 1999.




Figure 7: Pooled empirical forecast function. The figure shows the estimated forecast function if the data is pooled across all sectors using daily data. The lookback period of the input linear TSMOM signal is 21 day, 62 day or 260 days.

Table 1: Futures Contract Specifications. This table reports the list of the assets used in our analysis. Panel A describes the commodities, Panel B the rates and FX, and Panel C the equity indices

Ticker	Name	Exchange	Currency	Sector	Expiry Months
Panel A: Commodities					
LA1 Comdty	Aluminium	London Metal Exchange	USD	Base Metal	FGHJKMNQUVXZ
LP1 Comdty	Copper	London Metal Exchange	USD	Base Metal	FGHJKMNQUVXZ
LN1 Comdty	Nickel	London Metal Exchange	USD	Base Metal	FGHJKMNQUVXZ
LX1 Comdty	Zinc	London Metal Exchange	USD	Base Metal	FGHJKMNQUVXZ
CO1 Comdty	Brent	ICE Futures Europe Commodities	USD	Energy	FGHJKMNQUVXZ
QS1 Comdty	Gasoil	ICE Futures Europe Commodities	USD	Energy	FGHJKMNQUVXZ
KC1 Comdty	Coffee	ICE Futures US Softs	USD	Agriculturals	HKNUZ
CC1 Comdty	Cocoa	ICE Futures US Softs	USD	Agriculturals	HKNUZ
CT1 Comdty	Cotton	ICE Futures US Softs	USD	Agriculturals	HKNVZ
SB1 Comdty	Sugar	ICE Futures US Softs	USD	Agriculturals	HKNV
C 1 Comdty	Corn	Chicago Board of Trade	USD	Agriculturals	HKNUZ
S 1 Comdty	Soybean	Chicago Board of Trade	USD	Agriculturals	FHKNQUX
BO1 Comdty	SoybeanOil	Chicago Board of Trade	USD	Agriculturals	FHKNQUVZ
SM1 Comdty	SoybeanMeal	Chicago Board of Trade	USD	Agriculturals	FHKNQUVZ
W 1 Comdty	Wheat	Chicago Board of Trade	USD	Agriculturals	HKNUZ
LH1 Comdty	LeanHogs	Chicago Mercantile Exchange	USD	Agriculturals	GJKMNQVZ
LC1 Comdty	LiveCattle	Chicago Mercantile Exchange	USD	Agriculturals	GJMQVZ
CL1 Comdty	Crude	New York Mercantile Exchange	USD	Energy	FGHJKMNQUVXZ
XB1 Comdty	Gasoline	New York Mercantile Exchange	USD	Energy	FGHJKMNQUVXZ
HO1 Comdty	ULSD	New York Mercantile Exchange	USD	Energy	FGHJKMNQUVXZ
NG1 Comdty	NatGas	New York Mercantile Exchange	USD	Energy	FGHJKMNQUVXZ
GC1 Comdty	Gold	Commodity Exchange, Inc.	USD	Precious Metal	GJMQVZ
SI1 Comdty	Silver	Commodity Exchange, Inc.	USD	Precious Metal	HKNUZ
JA1 Comdty	Platinum	Osaka Exchange	JPY	Precious Metal	GJMQVZ
Panel B: Rates and FX					
YM1 Comdty	Australian3y	ASX Trade24	AUD	Rates	HMUZ
XM1 Comdty	Australian10y	ASX Trade24	AUD	Rates	HMUZ
DU1 Comdty	Schatz	Eurex	EUR	Rates	HMUZ
OE1 Comdty	Bobl	Eurex	EUR	Rates	HMUZ
RX1 Comdty	Bund	Eurex	EUR	Rates	HMUZ
UB1 Comdty	Buxl	Eurex	EUR	Rates	HMUZ
CN1 Comdty	Canadian10y	Montreal Exchange	CAD	Rates	HMUZ
JB1 Comdty	Japan10y	Osaka Exchange	JPY	Rates	HMUZ
G 1 Comdty	Gilt	ICE Futures Europe Financials	GBP	Rates	HMUZ
TU1 Comdty	US2y	Chicago Board of Trade	USD	Rates	HMUZ
FV1 Comdty	US5y	Chicago Board of Trade	USD	Rates	HMUZ
TY1 Comdty	US10y	Chicago Board of Trade	USD	Rates	HMUZ
US1 Comdty	US (long)	Chicago Board of Trade	USD	Rates	HMUZ
AD1 Curncy	AUD	Chicago Mercantile Exchange	USD	Currencies	HMUZ
BP1 Curncy	GBP	Chicago Mercantile Exchange	USD	Currencies	HMUZ
EC1 Curncy	EUR	Chicago Mercantile Exchange	USD	Currencies	HMUZ
JY1 Curncy	JPY	Chicago Mercantile Exchange	USD	Currencies	HMUZ
NO1 Curncy	NOK	Chicago Mercantile Exchange	USD	Currencies	HMUZ
SE1 Curncy	SEK	Chicago Mercantile Exchange	USD	Currencies	HMUZ
SF1 Curncy	CHF	Chicago Mercantile Exchange	USD	Currencies	HMUZ
CD1 Curncy	CAD	Chicago Mercantile Exchange	USD	Currencies	HMUZ

NV1 Curncy	NZD	Chicago Mercantile Exchange	USD	Currencies	HMUZ
------------	-----	-----------------------------	-----	------------	------

Panel C: Equities

XP1 Index	SPI	ASX Trade24	AUD	Equity Index	HMUZ
CF1 Index	CAC	Euronext Derivatives Paris	EUR	Equity Index	FGHJKMNQUVXZ
GX1 Index	DAX	Eurex	EUR	Equity Index	HMUZ
ST1 Index	MIB	Borsa Italiana (IDEM)	EUR	Equity Index	HMUZ
TP1 Index	TOPIX	Osaka Exchange	JPY	Equity Index	HMUZ
IB1 Index	IBEX	Meff Renta Variable (Madrid)	EUR	Equity Index	FGHJKMNQUVXZ
Z 1 Index	FTSE	ICE Futures Europe Financials	GBP	Equity Index	HMUZ
ES1 Index	SPX	Chicago Mercantile Exchange	USD	Equity Index	HMUZ

Table 2: Summary statistics. This table reports summary statistics of the assets used in our analysis. The time series start of each asset is reported in the last column.

Ticker	Returns (ann %)	Vol (ann %)	Skew	Kurt	Start date
Panel A: Commodities					
LA1 Comdty	-3.7	20.2	-0.2	2.8	1997-07-25
LP1 Comdty	4.6	24.8	-0.1	4.8	1997-07-01
LN1 Comdty	4.5	34.6	-0.1	3.9	1997-07-25
LX1 Comdty	-1.1	28.1	-0.2	3.4	1997-07-25
CO1 Comdty	5.9	35.2	-1.1	20.9	1988-06-27
QS1 Comdty	5.5	32.5	-0.7	14.9	1989-07-05
KC1 Comdty	-8.2	33.9	0.2	6.9	1980-01-02
CC1 Comdty	-7.8	28.6	0.0	2.2	1980-01-02
CT1 Comdty	-2.6	23.3	-0.0	1.7	1980-01-02
SB1 Comdty	-5.5	36.3	-0.2	3.7	1980-01-02
C 1 Comdty	-7.4	22.6	-0.0	3.0	1980-01-02
S 1 Comdty	0.5	21.5	-0.2	2.5	1980-01-02
BO1 Comdty	-4.4	22.4	0.1	1.8	1980-01-02
SM1 Comdty	4.8	23.5	-0.0	2.4	1980-01-02
W 1 Comdty	-8.6	25.5	0.1	2.5	1980-01-02
LH1 Comdty	-0.4	24.1	-0.2	2.7	1986-04-03
LC1 Comdty	1.7	15.1	-0.1	1.6	1980-01-02
CL1 Comdty	1.1	38.3	-2.1	51.6	1983-04-01
XB1 Comdty	1.6	40.1	-1.4	26.6	2005-10-05
HO1 Comdty	6.2	34.2	-0.8	16.7	1986-07-02
NG1 Comdty	-19.6	48.0	0.0	2.9	1990-04-05
GC1 Comdty	-1.7	18.3	-0.2	7.7	1980-01-02
SI1 Comdty	-5.3	29.8	-0.6	5.9	1980-01-02
JA1 Comdty	0.5	22.3	1.9	65.6	1984-01-30
Panel B: Rates and FX					
YM1 Comdty	0.7	1.2	-0.1	5.3	1989-12-13
XM1 Comdty	0.5	1.2	-0.3	5.2	1987-09-22
DU1 Comdty	0.7	1.2	-0.3	6.3	1997-03-11
OE1 Comdty	2.6	3.1	-0.2	2.5	1991-10-08

RX1 Comdty	4.0	5.2	-0.2	2.0	1990-11-27
UB1 Comdty	5.2	11.1	-0.2	2.7	1998-10-06
CN1 Comdty	3.5	5.9	-0.2	2.2	1989-09-19
JB1 Comdty	2.7	4.5	-0.7	13.4	1985-10-22
G 1 Comdty	2.6	7.2	-0.0	3.7	1982-11-22
TU1 Comdty	1.3	1.5	-0.0	6.3	1990-06-27
FV1 Comdty	2.6	3.8	-0.1	3.4	1988-05-24
TY1 Comdty	4.2	6.4	0.1	3.7	1982-05-05
US1 Comdty	4.1	11.0	-0.1	2.5	1980-01-02
AD1 Curncy	2.7	11.5	-0.6	10.1	1987-01-14
BP1 Curncy	0.9	9.7	-0.5	7.0	1986-05-29
EC1 Curncy	-0.3	9.4	-0.0	1.7	1998-05-21
JY1 Curncy	-0.9	10.6	0.5	6.8	1986-05-26
NO1 Curncy	0.4	12.3	-0.5	5.3	2002-05-20
SE1 Curncy	0.7	11.7	0.0	4.8	2002-05-20
SF1 Curncy	0.7	11.3	1.1	30.9	1986-04-08
CD1 Curncy	0.8	7.4	-0.1	6.2	1986-04-07
NV1 Curncy	2.3	12.6	-0.3	3.7	1997-05-09

Panel C: Equities

XP1 Index	3.8	16.5	-0.7	9.1	2000-05-04
CF1 Index	3.8	21.7	-0.2	5.8	1988-12-09
GX1 Index	4.9	22.2	-0.2	6.5	1990-11-27
ST1 Index	1.8	23.6	-0.7	9.8	2004-03-24
TP1 Index	-0.4	22.3	-0.1	8.3	1990-05-18
IB1 Index	5.0	23.1	-0.3	6.7	1992-07-22
Z 1 Index	3.0	17.8	-0.2	6.5	1988-03-01
ES1 Index	5.3	19.6	-0.3	11.7	1997-09-11

Table 3: Sharpe ratio and Out-of-sample R^2 - daily strategy. Sharpe ratios for the different TSMOM signals are reported in the first column of each panel. The remaining columns display p -values for pairwise tests of Sharpe ratio differences between signals. For instance, in Panel B, the value 0.08 in the "linear" row and second column corresponds to the p -value from testing whether the Sharpe ratio of the linear forecast (0.33) differs from that of the nonlinear forecast (0.47). Each panel corresponds to a different lookback period: Panel A uses a 1-month lookback, Panel B a 3-month lookback, and Panel C a 12-month lookback. Daily data.

Panel A: Look-back 1 month					
	SR	p -values			OOS- R^2
	(1)	(2)	(3)	(4)	(5)
Linear	0.24	(0.00)	(0.00)	(0.00)	0.08
F&S (nonlinear)	0.59***	-	(0.02)	(0.95)	0.11
Empirical nonlinear	0.83***	-	-	(0.98)	0.10
Binary	0.53***	-	-	-	0.11

Panel B: Look-back 3 months					
	SR	p -values			OOS- R^2
	(1)	(2)	(3)	(4)	(5)
Linear	0.33*	(0.08)	(0.04)	(0.11)	0.09
F&S (nonlinear)	0.47***	-	(0.01)	(0.74)	0.12
Empirical nonlinear	0.55***	-	-	(0.96)	0.11
Binary	0.45***	-	-	-	0.10

Panel C: Look-back 12 months					
	SR	p -values			OOS- R^2
	(1)	(2)	(3)	(4)	(5)
Linear	0.70***	(0.03)	(0.01)	(0.01)	0.09
F&S (nonlinear)	0.83***	-	(0.34)	(0.07)	0.10
Empirical nonlinear	0.84***	-	-	(0.11)	0.10
Binary	0.88***	-	-	-	0.10

Table 4: Optimal weight to TSMOM. This table reports the ex-post optimal weight to TSMOM for investors trading risk parity strategies, i.e. $w_{opt} = \max \text{Sharpe}((1-w)r_R + w r_{TSMOM})$. Each column displays the results for a different look-back period, e.g., 1 month, 3 months, and 12 months. Daily data.

	Look-back horizon		
	1 month	3 months	12 months
Linear	0.33	0.39	0.61
F&S (nonlinear)	0.58	0.54	0.67
Empirical nonlinear	0.70	0.59	0.68
Binary	0.57	0.54	0.70

Table 5: Performance of TSMOM strategies using a single pooled estimated nonlinear function (daily). This table reports the Sharpe ratios of the various estimations in columns (1)-(4). Column (5) reports the p-value of the difference between the empirical nonlinear and linear Sharpe ratios. Statistical significance is denoted by *** for p-value < 0.01, ** for p-value < 0.05 and * for p-value < 0.10.

	Sharpe ratio				<i>p-value</i>
	<i>Linear</i>	<i>F&S</i>	<i>Empirical nonlinear</i>	<i>Binary</i>	
	(1)	(2)	(3)	(4)	(5)
Panel A: Look-back 1 month					
Commodities	0.51***	0.75***	0.76***	0.75***	0.02
Equity	-0.11	0.04	0.13	-0.02	0.17
Rates	0.12	0.39**	0.62***	0.31*	0.02
Panel B: Look-back 3 months					
Commodities	0.53***	0.56***	0.57***	0.59***	0.37
Equity	-0.11	0.08	0.17	0.03	0.04
Rates	0.18	0.30*	0.35**	0.27*	0.10
Panel C: Look-back 12 months					
Commodities	0.65***	0.71***	0.70***	0.75***	0.19
Equity	0.19	0.31*	0.30*	0.30*	0.05
Rates	0.52***	0.65***	0.67***	0.68***	0.03

Table 6: Sharpe ratio and Out-of-sample R^2 - monthly strategy. Sharpe ratios for the different TSMOM signals are reported in the first column of each panel. The remaining columns display p -values for pairwise tests of Sharpe ratio differences between signals. For instance, in Panel C, the value 0.09 in the "linear" row and second column corresponds to the p -value from testing whether the Sharpe ratio of the linear forecast (0.55) differs from that of the nonlinear forecast (0.65). Each panel corresponds to a different lookback period: Panel A uses a 1-month lookback, Panel B a 3-month lookback, and Panel C a 12-month lookback. Monthly data.

Panel A: Look-back 1 month					
	<i>SR</i>	<i>p-values</i>			<i>OOS-R^2</i>
	(1)	(2)	(3)	(4)	(5)
Linear	0.38**	(0.95)	(0.87)	(0.92)	0.76
F&S (nonlinear)	0.28*	-	(0.04)	(0.61)	1.40
Empirical nonlinear	0.32**	-	-	(0.88)	1.61
Binary	0.27*	-	-	-	0.68

Panel B: Look-back 3 months					
	<i>SR</i>	<i>p-values</i>			<i>OOS-R^2</i>
	(1)	(2)	(3)	(4)	(5)
Linear	0.46**	(0.65)	(0.78)	(0.74)	0.63
F&S (nonlinear)	0.43**	-	(0.85)	(0.66)	1.28
Empirical nonlinear	0.36**	-	-	(0.24)	1.49
Binary	0.41**	-	-	-	0.67

Panel C: Look-back 12 months					
	<i>SR</i>	<i>p-values</i>			<i>OOS-R^2</i>
	(1)	(2)	(3)	(4)	(5)
Linear	0.55***	(0.09)	(0.06)	(0.19)	0.37
F&S (nonlinear)	0.65***	-	(0.53)	(0.90)	1.10
Empirical nonlinear	0.65***	-	-	(0.87)	1.34
Binary	0.61***	-	-	-	0.50

Table 7: Performance of TSMOM forecasts using a single pooled estimated nonlinear function (monthly).
This table reports the Sharpe ratios of the various estimations in columns (1)-(4). Column (5) reports the p-value of the difference between the empirical non linear and linear Sharpe ratios. Statistical significance is denoted by *** for p-value < 0.01, ** for p-value < 0.05 and * for p-value < 0.10.

	Sharpe ratio				<i>p-value</i>
	<i>Linear</i>	<i>F&S</i>	<i>Empirical nonlinear</i>	<i>Binary</i>	
	(1)	(2)	(3)	(4)	(5)
Panel A: Look-back 1 month					
Commodities	0.44**	0.41**	0.42**	0.43**	0.63
Equity	0.09	-0.01	0.09	-0.09	0.37
Rates	0.26*	0.16	0.18	0.18	0.90
Panel B: Look-back 3 months					
Commodities	0.41**	0.39**	0.21	0.42**	0.84
Equity	0.25*	0.33**	0.36**	0.28*	0.01
Rates	0.33**	0.25*	0.21	0.23	0.90
Panel C : Look-back 12 months					
Commodities	0.45**	0.44***	0.46***	0.40**	0.44
Equity	0.31*	0.40**	0.41**	0.36**	0.10
Rates	0.36**	0.48***	0.46***	0.48***	0.10

Table 8: Sharpe ratio and Out-of-sample R^2 - HOLDING PERIOD: 1 MONTH. Panel A reports unconditional results, while Panels B–D present results conditional on different market states. We classify months with returns above 1.33% as high, below -1.33% as low, and within $\pm 1.33\%$ as medium, a threshold implying roughly equal observations across buckets. The columns (2)–(4) display p -values for pairwise tests of Sharpe ratio differences between signals. For instance, in Panel A, the value 0.09 in the "linear" row and second column corresponds to the p -value from testing whether the Sharpe ratio of the linear forecast (0.55) differs from that of the nonlinear forecast (0.65). Monthly data.

Look-back 12 months					
	<i>SR</i>	<i>p-values</i>		<i>OOS-R^2</i>	
	(1)	(2)	(3)	(4)	(5)
Panel A: Unconditional					
Linear	0.55***	(0.09)	(0.07)	(0.20)	0.37
F&S (nonlinear)	0.65***	nan	(0.57)	(0.89)	1.10
Emp. nonlinear	0.65***	nan	nan	(0.85)	1.34
Binary	0.61***	nan	nan	nan	0.50
Panel B: High					
Linear	0.15***	(0.82)	(0.74)	(0.90)	0.37
F&S (nonlinear)	0.06***	nan	(0.08)	(0.78)	1.10
Emp. nonlinear	0.10***	nan	nan	(0.94)	1.34
Binary	0.03***	nan	nan	nan	0.50
Panel C: Medium					
Linear	0.82***	(0.18)	(0.10)	(0.39)	0.37
F&S (nonlinear)	1.06***	nan	(0.09)	(0.99)	1.10
Emp. nonlinear	1.12***	nan	nan	(1.00)	1.34
Binary	0.90***	nan	nan	nan	0.50
Panel D: Low					
Linear	0.87***	(0.00)	(0.00)	(0.00)	0.37
F&S (nonlinear)	1.21***	nan	(1.00)	(0.46)	1.10
Emp. nonlinear	1.10***	nan	nan	(0.05)	1.34
Binary	1.21***	nan	nan	nan	0.50

Table 9: Sharpe ratio and Out-of-sample R^2 - HOLDING PERIOD: 3 MONTHS. Panel A reports unconditional results, while Panels B–D present results conditional on different market states. We classify months with returns above 1.33% as high, below -1.33% as low, and within $\pm 1.33\%$ as medium, a threshold implying roughly equal observations across buckets. The columns (2)–(4) display p -values for pairwise tests of Sharpe ratio differences between signals. For instance, in Panel A, the value 0.03 in the "linear" row and second column corresponds to the p -value from testing whether the Sharpe ratio of the linear forecast (0.83) differs from that of the nonlinear forecast (0.98). Monthly data.

Look-back 12 months					
	<i>SR</i>	<i>p-values</i>		<i>OOS-R^2</i>	
	(1)	(2)	(3)	(4)	(5)
Panel A: Unconditional					
Linear	0.83***	(0.03)	(0.01)	(0.09)	0.37
F&S (nonlinear)	0.98***	nan	(0.12)	(0.95)	1.10
Emp. nonlinear	1.03***	nan	nan	(0.97)	1.46
Binary	0.93***	nan	nan	nan	0.50
Panel B: High					
Linear	0.73***	(0.47)	(0.17)	(0.52)	0.37
F&S (nonlinear)	0.74***	nan	(0.11)	(0.72)	1.10
Emp. nonlinear	0.81***	nan	nan	(0.93)	1.46
Binary	0.72***	nan	nan	nan	0.50
Panel C: Medium					
Linear	0.71***	(0.26)	(0.06)	(0.52)	0.37
F&S (nonlinear)	0.86***	nan	(0.01)	(0.98)	1.10
Emp. nonlinear	1.10***	nan	nan	(1.00)	1.46
Binary	0.71***	nan	nan	nan	0.50
Panel D: Low					
Linear	1.05***	(0.00)	(0.02)	(0.00)	0.37
F&S (nonlinear)	1.40***	nan	(0.99)	(0.50)	1.10
Emp. nonlinear	1.25***	nan	nan	0.04	1.46
Binary	1.39***	nan	nan	nan	0.50

Table 10: Sharpe ratio and Out-of-sample R^2 - HOLDING PERIOD: 12 MONTHS. Panel A reports unconditional results, while Panels B–D present results conditional on different market states. We classify months with returns above 1.33% as high, below -1.33% as low, and within $\pm 1.33\%$ as medium, a threshold implying roughly equal observations across buckets. The columns (2)–(4) display p -values for pairwise tests of Sharpe ratio differences between signals. For instance, in Panel A, the value 0.29 in the "linear" row and second column corresponds to the p -value from testing whether the Sharpe ratio of the linear forecast (1.31) differs from that of the nonlinear forecast (1.37). Monthly data.

Look-back 12 months					
	<i>SR</i>	<i>p-values</i>		<i>OOS-R^2</i>	
	(1)	(2)	(3)	(4)	(5)
Panel A: Unconditional					
Linear	1.31***	(0.29)	(0.95)	(0.63)	0.37
F&S (nonlinear)	1.37***	nan	(0.99)	(1.00)	1.10
Emp. nonlinear	0.95***	nan	nan	(0.03)	1.55
Binary	1.28***	nan	nan	nan	0.50
Panel B: High					
Linear	1.35***	(0.89)	(0.98)	(0.94)	0.37
F&S (nonlinear)	1.16***	nan	(0.98)	(0.98)	1.10
Emp. nonlinear	0.69***	nan	nan	(0.05)	1.55
Binary	1.07***	nan	nan	nan	0.50
Panel C: Medium					
Linear	1.43***	(0.06)	(0.51)	(0.14)	0.37
F&S (nonlinear)	1.79***	nan	(0.87)	(0.93)	1.10
Emp. nonlinear	1.43***	nan	nan	(0.26)	1.55
Binary	1.68***	nan	nan	nan	0.50
Panel D: Low					
Linear	1.19***	(0.12)	(0.79)	(0.20)	0.37
F&S (nonlinear)	1.42***	nan	(0.97)	(0.83)	1.10
Emp. nonlinear	0.88***	nan	nan	(0.08)	1.55
Binary	1.33***	nan	nan	nan	0.50

Table 11: Alphas. Monthly frequency estimation.

Panel A: Forecasting/Holding Period: 1 month				
Lookback	unconditional	mkt > 1.33%	abs(mkt) <= 1.33%	mkt < -1.33%
1m	2.84***	1.82*	2.25**	1.63*
3m	3.80***	2.74***	2.19**	2.47***
12m	4.36***	3.65***	3.94***	3.30***
36m	-0.32	-0.56	-0.63	-0.50

Panel B: Forecasting/Holding Period: 3 months				
Lookback	unconditional	mkt > 1.33%	abs(mkt) <= 1.33%	mkt < -1.33%
1m	3.74***	2.67**	2.64**	2.58**
3m	4.87***	3.43***	3.19***	3.15***
12m	5.89***	5.38***	5.98***	4.84***
36m	0.26	-0.18	-0.42	-0.16

Panel C: Forecasting/Holding Period: 12 months				
Lookback	unconditional	mkt > 1.33%	abs(mkt) <= 1.33%	mkt < -1.33%
1m	4.03**	3.90***	3.56***	3.83***
3m	6.01***	4.38***	5.32***	3.89***
12m	8.57***	7.54***	5.35***	6.67***
36m	-3.35	-2.61*	-1.10	-2.44*

Table 12: Sharpe ratio and Out-of-sample R^2 - daily strategy for Fama-French factors. Sharpe ratios for the different TSMOM signals are reported in the first column of each panel. The remaining columns display p -values for pairwise tests of Sharpe ratio differences between signals. Each panel corresponds to a different lookback period: Panel A uses a 1-month lookback, Panel B a 3-month lookback, and Panel C a 12-month lookback. Daily data.

Panel A: Look-back 1 month					
	<i>SR</i>	<i>p-values</i>			<i>OOS-R^2</i>
	(1)	(2)	(3)	(4)	(5)
Linear	0.63***	(0.04)	(0.02)	(0.01)	-0.01
F&S (nonlinear)	0.87***	-	(0.80)	(0.32)	-0.10
Empirical nonlinear	0.82***	-	-	(0.12)	-0.08
Binary	0.90***	-	-	-	0.03

Panel B: Look-back 3 months					
	<i>SR</i>	<i>p-values</i>			<i>OOS-R^2</i>
	(1)	(2)	(3)	(4)	(5)
Linear	0.55***	(0.03)	(0.02)	(0.04)	0.03
F&S (nonlinear)	0.80***	-	(0.96)	(0.79)	-0.06
Empirical nonlinear	0.62***	-	-	(0.10)	0.02
Binary	0.75***	-	-	-	0.04

Panel C: Look-back 12 months					
	<i>SR</i>	<i>p-values</i>			<i>OOS-R^2</i>
	(1)	(2)	(3)	(4)	(5)
Linear	0.57***	(0.50)	(0.32)	(0.62)	0.01
F&S (nonlinear)	0.57***	-	(0.27)	(0.70)	0.02
Empirical nonlinear	0.60***	-	-	(0.80)	0.08
Binary	0.54***	-	-	-	0.01

Table 13: Sharpe ratio and Out-of-sample R^2 - monthly strategy for Fama-French factors. Sharpe ratios for the different TSMOM signals are reported in the first column of each panel. The remaining columns display p -values for pairwise tests of Sharpe ratio differences between signals. Each panel corresponds to a different lookback period: Panel A uses a 1-month lookback, Panel B a 3-month lookback, and Panel C a 12-month lookback. Monthly data.

Panel A: Look-back 1 month					
	<i>SR</i>	<i>p-values</i>			<i>OOS-R^2</i>
	(1)	(2)	(3)	(4)	(5)
Linear	0.47**	(0.50)	(0.41)	(0.91)	1.27
F&S (nonlinear)	0.49***	-	(0.36)	(1.00)	1.09
Empirical nonlinear	0.50***	-	-	(1.00)	0.87
Binary	0.28*	-	-	-	0.35

Panel B: Look-back 3 months					
	<i>SR</i>	<i>p-values</i>			<i>OOS-R^2</i>
	(1)	(2)	(3)	(4)	(5)
Linear	0.38**	(0.10)	(0.05)	(0.18)	0.07
F&S (nonlinear)	0.51***	-	(0.81)	(0.71)	0.95
Empirical nonlinear	0.45***	-	-	(0.36)	0.99
Binary	0.48***	-	-	-	0.56

Panel C: Look-back 12 months					
	<i>SR</i>	<i>p-values</i>			<i>OOS-R^2</i>
	(1)	(2)	(3)	(4)	(5)
Linear	0.31*	(0.10)	(0.08)	(0.31)	-0.65
F&S (nonlinear)	0.48***	-	(0.74)	(0.96)	0.87
Empirical nonlinear	0.46**	-	-	(0.92)	1.06
Binary	0.38**	-	-	-	0.29

Table 14: Sharpe Ratios and Appraisal Ratios of Time-Series and Cross-Sectional Momentum Strategies. HOLDING PERIOD: 1 Month. Panel A reports Sharpe ratios for time-series momentum (TSMOM) strategies and for cross-sectional momentum (XSMOM) at different lookback periods (1M, 3M, 12M, 36M). Panel B reports appraisal ratios of each TSMOM strategy relative to the corresponding XSMOM benchmark at the same lookback horizon. Results highlight that both linear and nonlinear TSMOM strategies significantly outperform XSMOM at the 3M and 12M horizons. The performance improvements can be decomposed into contributions from volatility scaling and from mean forecasting, with the latter delivering substantial incremental gains.

	1M	3M	12M	36M
Panel A: Sharpe Ratios				
Linear	0.38**	0.46**	0.55***	-0.02
F&S (nonlinear)	0.28*	0.43**	0.65***	-0.06
Emp. nonlinear	0.32*	0.36**	0.65***	-0.05
Binary	0.27*	0.41**	0.61***	-0.06
XSMOM	0.30*	0.25	0.48***	0.27*
Panel B: Appraisal Ratios (relative to XSMOM)				
Linear	0.23	0.42**	0.27*	-0.23
F&S (nonlinear)	0.09	0.37**	0.44***	-0.25
Emp. nonlinear	0.16	0.26*	0.44**	-0.02
Binary	0.09	0.35**	0.38**	-0.25

Notes: Sharpe ratios and appraisal ratios are reported with significance levels denoted by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Appraisal ratios in Panel B are computed relative to cross-sectional momentum (XSMOM) at the corresponding lookback horizon.

Table 15: Sharpe Ratios and Appraisal Ratios of Time-Series and Cross-Sectional Momentum Strategies. HOLDING PERIOD: 3 Months. Panel A reports Sharpe ratios for time-series momentum (TSMOM) strategies and for cross-sectional momentum (XSMOM) at different lookback periods (1M, 3M, 12M, 36M). Panel B reports appraisal ratios of each TSMOM strategy relative to the corresponding XSMOM benchmark at the same lookback horizon. Results highlight that both linear and nonlinear TSMOM strategies significantly outperform XSMOM at the 12M horizon. The performance improvements can again be decomposed into contributions from volatility scaling and from mean forecasting, with the latter delivering substantial incremental gains.

	1M	3M	12M	36M
Panel A: Sharpe Ratios				
Linear	0.47***	0.59***	0.83***	0.06
F&S (nonlinear)	0.41**	0.52***	0.98***	0.01
Emp. nonlinear	0.40**	0.48***	1.03***	-0.01
Binary	0.43**	0.52***	0.93***	0.01
XSMOM	0.36**	0.46***	0.73***	0.45**
Panel B: Appraisal Ratios (relative to XSMOM)				
Linear	0.31*	0.37**	0.40**	-0.34
F&S (nonlinear)	0.21	0.25	0.66***	-0.36
Emp. nonlinear	0.19	0.20	0.74***	-0.15
Binary	0.24*	0.25	0.58***	-0.36

Notes: Sharpe ratios and appraisal ratios are reported with significance levels denoted by stars: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. Appraisal ratios in Panel B are computed relative to cross-sectional momentum (XSMOM) at the corresponding lookback horizon.

	simple	linear / BSC	variance	dyn. reg.	non linear strategies		
					F&S	emp. nonl.	binary
simple	-	0.45**	0.24	0.09	0.62***	0.63***	0.56***
linear / BSC	-0.20	-	-0.05	-0.10	0.38**	0.41**	0.28*
variance	0.24	0.41**	-	0.12	0.55***	0.54***	0.50***
dyn. reg.	0.28	0.50**	0.28*	-	0.60***	0.60***	0.56***
nonl. FS	-0.31	-0.15	-0.08	-0.06	-	-0.00	-0.22
emp. nonl.	-0.34	-0.22	-0.08	-0.08	0.06	-	-0.16
binary	-0.26	-0.08	-0.05	-0.06	0.31*	0.26	-
Sharpe ratio	0.37**	0.55***	0.36**	0.25	0.65***	0.65***	0.61***

Table 16: Appraisal ratios: comparison. This table reports the appraisal ratios for the various strategies. In each column, the rows show the appraisal ratio (annualized) of the column name over the row name. As an example, the last number in the first row (0.56**) is the appraisal ratio of the binary strategy over the simple linear TSMOM. Vice versa, the last number in the first column (-0.26) is the appraisal ratio of the simple linear TSMOM over the binary strategy. The statistical significance is based on a one-sided test (hence, all negative appraisal ratios are insignificant). The last row, separated by a line contains the Sharpe ratios of the individual strategies.

Internet Appendix A

Table A.1: Subsample analysis (daily). Sharpe ratios by strategy, lookback and sub-period. Statistical significance against Null of non-positive Sharpe ratio.

Panel A: Linear			
<i>Lookback period</i>	<i>Sharpe ratio</i>		
	2000-2007	2007-2015	2015-2024
1 month	0.30	0.25	0.21
3 months	0.16	0.59 **	0.17
12 months	1.04 ***	0.66 **	0.44
Panel B: Nonlinear			
<i>Lookback period</i>	<i>Sharpe ratio</i>		
	2000-2007	2007-2015	2015-2024
1 month	0.93 ***	0.46	0.44 *
3 months	0.36	0.72 **	0.29
12 months	1.23 ***	0.90 ***	0.42 *
Panel C: Empirical nonlinear			
<i>Lookback period</i>	<i>Sharpe ratio</i>		
	2000-2007	2007-2015	2015-2024
1 month	1.38 ***	0.70 **	0.50 *
3 months	0.45 *	0.82 **	0.35
12 months	1.27 ***	0.87 ***	0.43 *
Panel D: Binary			
<i>Lookback period</i>	<i>Sharpe ratio</i>		
	2000-2007	2007-2015	2015-2024
1 month	0.89 ***	0.42	0.34
3 months	0.31	0.73 **	0.27
12 months	1.25 ***	0.94 ***	0.5 *

Internet Appendix B

Table B.1: Subsample analysis (monthly). Sharpe ratios by strategy, lookback and sub-period. Statistical significance against Null of non-positive Sharpe.

Panel A: Linear			
<i>Lookback period</i>	<i>Sharpe ratio</i>		
	2000-2007	2007-2015	2015-2024
1 month	0.61 **	0.38	0.24
3 months	0.65 **	0.67 **	0.05
12 months	0.87 ***	0.54 *	0.23
Panel B: Nonlinear			
<i>Lookback period</i>	<i>Sharpe ratio</i>		
	2000-2007	2007-2015	2015-2024
1 month	0.55 *	0.22	0.16
3 months	0.73 **	0.72 **	-0.07
12 months	1.24 ***	0.72 **	0.07
Panel C: Empirical nonlinear			
<i>Lookback period</i>	<i>Sharpe ratio</i>		
	2000-2007	2007-2015	2015-2024
1 month	0.57 *	0.31	0.2
3 months	0.67 **	0.65 **	-0.13
12 months	1.16 ***	0.69 **	0.13
Panel D: Binary			
<i>Lookback period</i>	<i>Sharpe ratio</i>		
	2000-2007	2007-2015	2015-2024
1 month	0.58 *	0.26	0.07
3 months	0.76 **	0.66 **	-0.07
12 months	1.14 ***	0.66 **	0.09

Internet Appendix C

Table C.1: Sharpe ratio and Out-of-sample R^2 - weekly strategy. Sharpe ratios for the different TSMOM signals are reported in the first column of each panel. The remaining columns display p -values for pairwise tests of Sharpe ratio differences between signals. For instance, in Panel C, the value 0.02 in the "linear" row and second column corresponds to the p -value from testing whether the Sharpe ratio of the linear forecast (0.59) differs from that of the nonlinear forecast (0.75). Each panel corresponds to a different lookback period: Panel A uses a 1-month lookback, Panel B a 3-month lookback, and Panel C a 12-month lookback. Weekly data.

Panel A: Look-back 1 month					
	<i>SR</i>	<i>p-values</i>			<i>OOS-R^2</i>
	(1)	(2)	(3)	(4)	(5)
Linear	0.27*	(0.07)	(0.04)	(0.09)	0.27
F&S (nonlinear)	0.41**	-	(0.04)	(0.70)	0.40
Empirical nonlinear	0.50***	-	-	0.93	0.39
Binary	0.39**	-	-	-	0.30

Panel B: Look-back 3 months					
	<i>SR</i>	<i>p-values</i>			<i>OOS-R^2</i>
	(1)	(2)	(3)	(4)	(5)
Linear	0.28*	(0.43)	(0.68)	(0.50)	0.25
F&S (nonlinear)	0.30*	-	(0.94)	(0.69)	0.40
Empirical nonlinear	0.23	-	-	(0.17)	0.47
Binary	0.29*	-	-	-	0.28

Panel C: Look-back 12 months					
	<i>SR</i>	<i>p-values</i>			<i>OOS-R^2</i>
	(1)	(2)	(3)	(4)	(5)
Linear	0.59***	(0.02)	(0.06)	(0.01)	0.24
F&S (nonlinear)	0.75***	-	(0.97)	(0.44)	0.29
Empirical nonlinear	0.66***	-	-	(0.02)	0.35
Binary	0.76***	-	-	-	0.28

Table C.2: Performance of TSMOM forecasts using a single pooled estimated nonlinear function (weekly). This table reports the Sharpe ratios of the various estimations in columns (1)-(4). Column (5) reports the p-value of the difference between the empirical non linear and linear Sharpe ratios. Statistical significance is denoted by *** for p-value < 0.01, ** for p-value < 0.05 and * for p-value < 0.10.

	Sharpe ratio				
	<i>Linear</i>	<i>F&S</i>	<i>Empirical nonlinear</i>	<i>Binary</i>	<i>p-value</i>
	(1)	(2)	(3)	(4)	(5)
Panel A: Look-back 1 month					
Commodities	0.50***	0.50***	0.53***	0.52***	0.33
Equity	-0.12	-0.03	0.09	-0.04	0.17
Rates	0.14	0.30*	0.33*	0.26	0.03
Panel B: Look-back 3 months					
Commodities	0.47***	0.44**	0.30*	0.41**	0.82
Equity	-0.12	-0.05	-0.03	-0.09	0.06
Rates	0.14	0.21	0.20	0.23	0.21
Panel C : Look-back 12 months					
Commodities	0.58***	0.64***	0.64***	0.66***	0.20
Equity	0.18	0.28*	0.27*	0.24	0.07
Rates	0.42**	0.59***	0.46**	0.59***	0.04

Internet Appendix D

Table D.1: Sharpe Ratio and Out-of-Sample R^2 — Holding Period: 1 Month. Panel A reports unconditional results, while Panels B–D present results conditional on different market states (High, Medium, Low). We define market states using a 20–60–20 split, where the bottom and top states each contain 20% of the observations. Column (2) reports Sharpe ratios. Columns (3)–(5) display p -values (in parentheses) for pairwise tests of Sharpe ratio differences between forecasts. For instance, in Panel A, the value (0.10) in the “linear” row and “nonlinear” column corresponds to the p -value for testing whether the Sharpe ratio of the linear forecast (0.55) differs from that of the nonlinear forecast (0.65). The last two columns report appraisal ratios (relative to TSMOM) and out-of-sample R^2 . Monthly data.

	<i>SR</i>	<i>p-values</i>			TSMOM-Appraisal	OOS- R^2
		F&S (nonlinear)	Emp. nonlinear	Binary		
Panel A: Unconditional						
Linear	0.55***	(0.10)	(0.06)	(0.19)	-	0.37
F&S (nonlinear)	0.65***	-	(0.53)	(0.90)	0.38**	0.50
Emp. nonlinear	0.65***	-	-	(0.83)	0.41**	0.52
Binary	0.61***	-	-	-	0.28*	0.50
Panel B: High						
Linear	-0.33***	(0.90)	(0.91)	(0.96)	-	-2.43
F&S (nonlinear)	-0.54***	-	(0.22)	(0.78)	-	-3.13
Emp. nonlinear	-0.50***	-	-	(0.93)	-	-3.32
Binary	-0.58***	-	-	-	-	-2.56
Panel C: Medium						
Linear	0.54***	(0.10)	(0.04)	(0.29)	-	0.12
F&S (nonlinear)	0.69***	-	(0.14)	(0.96)	-	0.51
Emp. nonlinear	0.72***	-	-	(0.99)	-	0.57
Binary	0.61***	-	-	-	-	0.45
Panel D: Low						
Linear	1.40***	(0.00)	(0.01)	(0.01)	-	0.90
F&S (nonlinear)	1.71***	-	(1.00)	(0.36)	-	1.22
Emp. nonlinear	1.60***	-	-	(0.11)	-	1.26
Binary	1.74***	-	-	-	-	0.94

Table D.2: Sharpe Ratio and Out-of-Sample R^2 — Holding Period: 3 Months. Panel A reports unconditional results, while Panels B–D present results conditional on different market states (High, Medium, Low). We define market states using a 20–60–20 split, where the bottom and top states each contain 20% of the observations. Column (2) reports Sharpe ratios. Columns (3)–(5) display p -values (in parentheses) for pairwise tests of Sharpe ratio differences between forecasts. The last two columns report appraisal ratios (relative to TSMOM) and out-of-sample R^2 . Monthly data.

	SR	<i>p-values</i>			TSMOM-Appraisal	OOS- R^2
		F&S (nonlinear)	Emp. nonlinear	Binary		
Panel A: Unconditional						
Linear	0.83***	(0.04)	(0.00)	(0.08)	-	-0.50
F&S (nonlinear)	0.98***	-	(0.14)	(0.96)	0.57***	-0.38
Emp. nonlinear	1.03***	-	-	(0.96)	0.73***	-0.13
Binary	0.93***	-	-	-	0.46**	-0.26
Panel B: High						
Linear	0.12***	(0.38)	(0.29)	(0.53)	-	-6.00
F&S (nonlinear)	0.17***	-	(0.39)	(0.95)	-	-6.28
Emp. nonlinear	0.19***	-	-	(0.83)	-	-6.40
Binary	0.11***	-	-	-	-	-5.55
Panel C: Medium						
Linear	0.81***	(0.23)	(0.01)	(0.34)	-	-0.79
F&S (nonlinear)	0.91***	-	(0.01)	(0.93)	-	-1.31
Emp. nonlinear	1.06***	-	-	(1.00)	-	-0.98
Binary	0.85***	-	-	-	-	-1.14
Panel D: Low						
Linear	1.55***	(0.00)	(0.05)	(0.00)	-	0.85
F&S (nonlinear)	1.98***	-	(0.99)	(0.38)	-	2.42
Emp. nonlinear	1.75***	-	-	(0.02)	-	2.49
Binary	2.00***	-	-	-	-	2.30