

Demystifying Time-Series Momentum Strategies: Volatility Estimators, Trading Rules and Pairwise Correlations*

NICK BALTAS[†] AND ROBERT KOSOWSKI[‡]

October 1, 2015

ABSTRACT

Motivated by studies of the impact of frictions on asset prices, we examine the effect of key components of time-series momentum strategies on their turnover and performance from 1974 until 2013. We show that more efficient volatility estimation and price trend detection significantly reduce portfolio turnover and therefore rebalancing costs. The poor performance of time-series momentum strategies during the post-2008 period is explained by an increased level of pairwise correlations. We propose a novel correlation-based leverage-adjustment to the strategy's weighting scheme and show that it improves performance by safeguarding against tail risk, even after accounting for realistic transaction costs.

JEL CLASSIFICATION CODES: E37, G11, G15, F37

KEY WORDS: Time-series momentum; Constant-volatility; Trading rule; Pairwise correlations; Turnover.

*Comments by Yoav Git, Mark Hutchinson, Nadia Linciano, Pedro Saffi, Mark Salmon, Stephen Satchell, Laurens Swinkels and participants at the UBS Annual Quantitative Conference (April 2013), the IV World Finance Conference (July 2013), the 67th European Meeting of the Econometric Society (Aug. 2013), the FMA Hedge Fund Consortium (Nov. 2014) and the FMA European Conference (June 2015) are gratefully acknowledged. Further comments are warmly welcomed, including references to related papers that have been inadvertently overlooked. Financial support from INQUIRE Europe is gratefully acknowledged. The views expressed in this article are those of the authors only and no other representation to INQUIRE Europe or UBS Investment Bank should be attributed. The paper has been previously circulated with the title "Improving Time-Series Momentum Strategies: The Role of Volatility Estimators and Trading Signals".

[†]Corresponding Author; (i) UBS Investment Bank, EC2M 2PP, London, UK, (ii) Imperial College Business School, South Kensington Campus, SW7 2AZ, London, UK, (iii) Queen Mary University of London, Mile End Road, E1 4NS, London, UK; e-mail: n.baltas@imperial.ac.uk.

[‡](i) Imperial College Business School, South Kensington Campus, SW7 2AZ, London, UK, (ii) Oxford-Man Institute of Quantitative Finance, University of Oxford, Walton Well Road, OX2 6ED, Oxford, UK, (iii) CEPR; e-mail: r.kosowski@imperial.ac.uk.

1. Introduction

Managed futures funds, also known as Commodity Trading Advisors (CTAs), constitute a significant part of the hedge fund industry. Using BarclayHedge estimates at the end of 2014, managed futures funds manage a total of \$318bn. of assets, which is about 11% of the \$2.8tr. hedge fund industry. These funds typically trade futures contracts on assets in various asset classes (equity indices, commodities, government bonds and FX rates) and profit from systematic price trends¹ by the means of time-series momentum strategies; Moskowitz, Ooi and Pedersen (2012) are the first to comprehensively study these strategies in the academic literature, whereas Hurst, Ooi and Pedersen (2013) and Baltas and Kosowski (2013) provide statistical evidence that managed futures and CTA funds do employ such strategies.²

Time-series momentum strategies are constructed using long and short positions based on a simple momentum-based trading rule, which is the sign of the past return. The weighting scheme that is employed is an inverse-volatility scheme (Moskowitz et al. 2012); the correlation structure of the portfolio constituents is not typically incorporated in the weighting scheme. These strategies have recently received considerable investor attention, because, on the one hand they provided impressive diversification benefits during the recent global financial crisis (GFC) of 2008, but on the other hand, they have exhibited rather poor performance during the subsequent post-crisis period between 2009 and 2013. One of the reasons that has been claimed to be responsible for this recent underperformance has been the increased level of correlations across markets and asset classes in the post-2008 period (Baltas and Kosowski 2013).

The objective of this paper is twofold. First, we focus on the portfolio turnover implications of the two key inputs in time-series momentum strategies, namely the volatility estimator that is used to scale the asset positions and the trading rule that is used to identify the price trends. In particular, we explore the benefits in terms of turnover reduction from employing more efficient volatility estimates and from identifying more accurately the strength of price trends by means of alternative trading rules. Second,

¹Using BarclayHedge estimates at the end of 2014, systematic traders manage \$293bn. of assets, which is about 92% of the \$318bn. of assets of the managed futures industry.

²Time-series momentum strategies are structurally different from the cross-sectional momentum strategies, “winners minus losers”, that have been heavily studied in the academic literature, starting from the works of Jegadeesh and Titman (1993, 2001) in the equity markets. Cross-sectional momentum effects in futures markets across multiple asset classes have been recently documented by Pirrong (2005), Miffre and Rallis (2007) and Asness, Moskowitz and Pedersen (2013). The time-series momentum strategies rely heavily on return serial correlation patterns, whereas the cross-sectional momentum strategies are a long-short zero-cost portfolios of securities with the best and worst relative performance during a lookback period, typically equal to 12 months.

the paper studies the dependence of the performance of time-series momentum strategies on the level of pairwise correlations, with special attention paid to their recent underperformance following the GFC. We introduce a correlation adjustment, which is applied at the portfolio level and reduces the employed leverage at periods of high co-movement, in order to safeguard against tail risk.

As part of our turnover reduction analysis, we first document the economic value of using a volatility estimator with desirable theoretical properties. In particular, in the spirit of Fleming, Kirby and Ostdiek (2003), we hypothesise that more efficient and accurate estimators, than those constructed using daily close-to-close returns, can safeguard against excessive rebalancing and therefore reduce the turnover and improve the performance of the strategy after transaction costs. By employing a range-based estimator (constructed using open-high-low-close prices), such as the one proposed by Yang and Zhang (2000), we empirically find that the turnover of the strategy is reduced by almost one tenth without affecting the risk-adjusted performance of the strategy before incorporating any trading costs. Importantly, the benefit in the turnover reduction is not due to a small number of portfolio constituents, but instead it is found to be pervasive across all portfolio constituents from all asset classes.

The typical momentum trading rule in the literature uses the sign of the past return over the past 12 months (Moskowitz et al. 2012). Our analysis shows that the frequency at which a trading rule switches between long and short positions can dramatically affect the portfolio turnover. Intuitively, avoiding the excessive position changes when no significant price trend exists can significantly reduce the turnover and therefore improve the performance of the strategy after accounting for transaction costs. Therefore, we evaluate the properties and investment implications of a trading rule that only instructs taking a long or a short position when the underlying price trend is statistically significant. By effectively timing the entry to and the exit from a particular position we find that the turnover of the strategy falls largely by two thirds compared to using the sign of the past returns.

Our second objective is to study the dependence of the performance of the time-series momentum strategy on the level of pairwise correlations of portfolio constituents. This analysis is motivated by the findings in Baltas and Kosowski (2013), who, apart from documenting the business cycle performance of the strategy, also highlight its poor performance after 2008. The authors explain that the underperformance can be due to (i) capacity constraints in the futures markets, (ii) a lack of trends for each asset

or (iii) increased correlations across assets. They find no evidence of capacity constraints based on two different methodologies, but they do show that correlations between futures markets have experienced a significant increase in the period from 2008 to 2013.

We show that incorporating the pairwise correlations of the constituent assets into the weighting scheme of the time-series momentum strategy not only sheds light on the return drivers of time-series momentum strategies, but can also significantly improve their out-of-sample performance in the years following the GFC. We investigate the interplay between the pairwise correlations of portfolio constituents and the portfolio volatility and extend the formulation of the standard time-series momentum strategy by introducing a correlation factor in the weighting scheme that increases (decreases) the leverage of portfolio constituents in periods of low (high) average pairwise correlation. This adjustment improves the risk-adjusted performance of the strategy, because it actively safeguards against crash risk. The improvement is relatively more pronounced over the most recent post-crisis period 2009-2013 during which pairwise correlations across assets and asset classes dramatically increased, thus, diminishing diversification benefits. In particular, the Sharpe ratio of the strategy over this period more than doubles (from 0.14 to 0.29) after employing the correlation adjustment. Nevertheless, the performance benefit does not come at no cost, as the turnover of the strategy increases.

The various methodological adjustments that we discuss in this paper have different turnover implications for the time-series momentum strategy; more efficient volatility estimation and price trend detection result in turnover reduction, whereas incorporating pairwise correlations increases the turnover at the benefit of improved portfolio diversification. In order to evaluate the economic impact of differences in turnover on strategy returns, we construct our time-series momentum strategies using a comprehensive dataset of 75 futures contracts over a period of 36 years and we use the simple approximation for transaction costs of Frazzini, Israel and Moskowitz (2012). Under this costs model, the realised costs of a strategy are proportional to portfolio turnover and an average market impact level; the market impact is assumed constant for our analysis, even though in practice the market impact typically depends, among other things, on the order size. Contrary to expectations and based on this model, the turnover and costs reduction that is achieved by more efficient volatility estimation and price trend detection does not lead to significantly higher risk-adjusted performance of the time-series momentum strategy. However, the performance improvement of the strategy due to the incorporation of pairwise correlation in the weighting

scheme is shown to be robust to the increased trading costs.

Our paper is related to several streams of the literature including work on the economic value of volatility-timing and the importance of volatility estimation efficiency on dynamic portfolio construction and turnover. Fleming, Kirby and Ostdiek (2001), Ilmanen and Kizer (2012), Kirby and Ostdiek (2012) and Hallerbach (2012) highlight the benefits of volatility-timing, while Barroso and Santa-Clara (2015) and Daniel and Moskowitz (2015) examine its effect on the performance of cross-sectional equity momentum strategies. Fleming et al. (2003) investigate the performance and turnover benefits for a mean-variance portfolio from using more efficient estimates of volatility. We make use of range volatility estimators in order to increase the estimation efficiency of daily estimators.

The rest of the paper is organised as follows. Section 2 provides an overview of the dataset and Section 3 describes the construction of the time-series momentum strategy, explores the dependence of the strategy's turnover on volatility estimator and trading rule and introduces the correlation adjustment. The empirical results of the effects of volatility estimator and trading rule on the turnover of time-series momentum strategies are presented in Section 4. Section 5 discusses the recent underperformance of the time-series momentum strategies and presents our empirical results on the effect of incorporating pairwise correlations in the weighting scheme onto the performance of these strategies. Section 6 reports the transaction cost implications of all the methodological alterations that are presented over the course of the paper and finally, Section 7 concludes.

2. Data Description

The dataset that we use is based on the one used by Baltas and Kosowski (2013) and consists of daily opening, high, low and closing futures prices for 75 assets across all asset classes: 26 commodities, 23 equity indices, 7 currencies and 19 short-term, medium-term and long-term government bonds; see Table I. It is obtained from Tick Data and the sample period is from December 1974 (not all contracts start in December 1974; Table I reports the starting month and year of each contract) to February 2013. Since the contracts of different assets are traded on various exchanges each with different trading hours and holidays, the data series are appropriately aligned by filling forward any missing prices. Finally and especially for equity indices, we also obtain spot prices from Datastream and backfill the respective

futures series for periods prior to the availability of futures data.³

Futures contracts are short-lived instruments and are only active for a few months until the delivery date. Additionally, entering a futures contract is, in theory, a free of cost investment and in practice only implies a small (relative to a spot transaction) initial margin payment, hence rendering futures highly levered investments. These features of futures contracts give rise to two key issues that we carefully address below, namely (a) the construction of single continuous price time-series per asset suitable for backtesting and (b) the calculation of holding period returns.

First, in order to construct a continuous series of futures prices for each asset, we appropriately splice together different contracts. Following the standard approach in the literature (e.g. de Roon et al. 2000, Miffre and Rallis 2007, Moskowitz et al. 2012), we use the most liquid futures contract at each point in time and we roll over contracts so that we always trade the most liquid contract. The most liquid contract is typically the nearest-to-delivery (“front”) contract up until a few days/weeks before delivery, when the second-to-delivery (“first-back”) contract becomes the most liquid one and a rollover takes place.

An important issue for the construction of continuous price series of a futures contract is the price adjustment on a roll date. The two contracts that participate in a rollover do not typically trade at the same price. If the time-series of these contracts were to be spliced together without any further adjustment, then an artificial non-traded return would appear on the rollover day, which would bias the mean return upwards or downwards for an asset that is on average in contango or backwardation respectively. For that purpose, we backwards ratio-adjust the futures series at each roll date, i.e. we multiply the entire history of the asset by the ratio of the prevailing futures prices of the new and the old contracts. Hence, the entire price history up to the roll date is scaled accordingly so that no artificial return exists in the continuous data series.⁴

Second, having obtained single price data series for each asset, we need to calculate daily excess

³de Roon, Nijman and Veld (2000) and Moskowitz et al. (2012) find that equity index returns calculated using spot price series or nearest-to-delivery futures series are largely correlated. In unreported results, we confirm this and find that our results remain qualitatively unchanged without the equity spot price backfill.

⁴Another price adjustment technique is to add/subtract to the entire history the level difference between the prevailing futures prices of the two contracts involved in a rollover (backwards difference adjustment). The disadvantage of this technique is that it distorts the historical returns as the price level changes in absolute terms. In fact, the historical returns are upwards or downwards biased for contracts that are on average in backwardation or contango respectively. Instead, backwards ratio adjustment only scales the price series, hence it leaves percentage changes unaffected and results in a tradable series that can be used for backtesting.

returns. As already mentioned, calculating futures holding period returns is not as straightforward as it is for spot transactions and requires additional assumptions regarding the margin payments. For that purpose, let $F_{t,T}$ and $F_{t+1,T}$ denote the prevailing futures prices of a futures contract with maturity T at the end of months t and $t + 1$ respectively. Additionally, assume that the contract is not within its delivery month, hence $t < t + 1 < T$. Entering a futures contract at time t implies an initial margin payment of M_t that earns the risk-free rate, r_t^f during the life of the contract. During the course of a month, assuming no variation margin payments, the margin account will have accumulated an amount equal to $M_t (1 + r_t^f) + (F_{t+1,T} - F_{t,T})$. Therefore, the holding period return for the futures contract in excess of the risk-free rate is:

$$r_{t,t+1}^{margin} = \frac{[M_t (1 + r_t^f) + (F_{t+1,T} - F_{t,T})] - M_t}{M_t} - r_t^f = \frac{F_{t+1,T} - F_{t,T}}{M_t} \quad (1)$$

If we assume that the initial margin requirement equals the prevailing futures price, i.e. $M_t = F_{t,T}$ then we can calculate the fully collateralised return in excess of the risk-free rate as follows:

$$r_{t,t+1} = \frac{F_{t+1,T} - F_{t,T}}{F_{t,T}} \quad (2)$$

Interestingly, the excess return calculation for a fully collateralised futures transaction takes the same form as a total return calculation for a cash equity spot transaction.

Using Equation (2), we construct daily excess close-to-close fully collateralised returns, which are then compounded to generate monthly returns.⁵ Table I presents summary monthly return statistics for all assets and asset classes. In line with the futures literature (e.g. see de Roon et al. 2000, Moskowitz et al. 2012), we find that there is large cross-sectional variation in the return distributions of the different assets. In total, 67 out of 75 futures contracts have a positive unconditional mean excess return, 29 of which statistically significant at the 10% level. Currency and commodity futures have insignificant mean returns with only few exceptions. All but four assets have leptokurtic return distributions (fat tails) and, as expected, almost all equity futures have negative skewness. More importantly, the cross-sectional variation in volatility is substantial. Commodity and equity futures exhibit the largest volatilities, followed

⁵This procedure is fairly standard in the futures literature. Among others, see Bessembinder (1992), Bessembinder (1993), Gorton and Rouwenhorst (2006), Gorton, Hayashi and Rouwenhorst (2007), Miffre and Rallis (2007), Pesaran, Schleicher and Zaffaroni (2009), Fuertes, Miffre and Rallis (2010) and Moskowitz et al. (2012).

by the currencies and lastly by the bond futures, which have the lowest volatilities in the cross-section.

[Table I about here]

3. Methodology

In this section we provide the background in the construction of time-series momentum strategies, which is necessary to understand why their performance is dependent on the efficiency of volatility estimation, the momentum trading rule and the level of pairwise correlation of its constituents. These strategies can be viewed as an extension of constant-volatility (volatility-targeting) strategies, which explains the dependence of strategies' turnover on the efficiency of the volatility estimation and on the nature of the trading rule. At the end of the section we extend the strategy design by explicitly incorporating the pairwise correlations of the constituents.

3.1. Constant-Volatility and Time-Series Momentum Strategies

In the previous section, we discussed the return construction of a fully collateralised futures position. In practice, the initial margin requirement is a fraction of the prevailing futures price and is typically a function of the historical risk profile of the underlying asset. If we therefore express the initial margin requirement as the product of the underlying asset's volatility and its futures price, i.e. $M_t = \sigma_t \cdot F_{t,T}$, we can deduce from Equation (1) a levered holding period return in excess of the risk-free rate as follows:

$$r_{t,t+1}^{lev} = \frac{F_{t+1,T} - F_{t,T}}{\sigma_t \cdot F_{t,T}} = \frac{1}{\sigma_t} \cdot r_{t,t+1} \quad (3)$$

It is worth noting that the above result can also be interpreted as a long-only constant-volatility strategy, with the target level of volatility being equal to 100%. Denoting by σ_{tgt} a desired level of target volatility, we can generalise the concept to a single-asset constant-volatility (*cvol*) strategy:

$$r_{t,t+1}^{cvol} = \frac{\sigma_{tgt}}{\sigma_t} \cdot r_{t,t+1} \quad (4)$$

The concept of constant-volatility (also known as volatility-targeting or volatility-timing) has been first

highlighted by Fleming et al. (2001, 2003) and more recently by Kirby and Ostdiek (2012), Ilmanen and Kizer (2012) and Hallerbach (2012). This series of papers documents that volatility-timing can result in desirable properties for the portfolio like lower turnover and larger Sharpe ratio.

A constant-volatility strategy across assets (to differentiate from the single-asset strategy, we hereafter use capital letters, CVOL) can be simply formed as the average (equally-weighted portfolio) of individual constant-volatility strategies:

$$r_{t,t+1}^{CVOL} = \frac{1}{N_t} \sum_{i=1}^{N_t} r_{t,t+1}^{i,cvol} \quad (5)$$

$$= \frac{1}{N_t} \sum_{i=1}^{N_t} \frac{\sigma_{tgt}}{\sigma_t^i} \cdot r_{t,t+1}^i \quad (6)$$

where N_t is the number of available assets at time t . The target volatility of each asset remains σ_{tgt} , however the volatility of the portfolio is expected to be relatively lower due to diversification. In fact, the volatility of the portfolio would only be equal to this upper bound of σ_{tgt} , if all the assets were perfectly correlated, which is not typically the case. Further details on the effect of pairwise correlations are presented later in this section.

A time-series momentum strategy (TSMOM, hereafter), also known as a trend-following strategy, is an extension to the long-only CVOL strategy of Equation (6) and involves both long and short positions. These are determined by each asset's recent performance over some lookback period, as captured by an appropriately designed trading rule denoted by X :

$$r_{t,t+1}^{TSMOM} = \frac{1}{N_t} \sum_{i=1}^{N_t} X_t^i \cdot r_{t,t+1}^{i,cvol} \quad (7)$$

$$= \frac{1}{N_t} \sum_{i=1}^{N_t} X_t^i \cdot \frac{\sigma_{tgt}}{\sigma_t^i} \cdot r_{t,t+1}^i \quad (8)$$

The nature of the trading rule is critical for the performance of the strategy. In its simplest form (as in Moskowitz et al. 2012, Hurst et al. 2013, Baltas and Kosowski 2013), the TSMOM strategy uses the sign of the past 12-month return to determine the type of position for each portfolio constituent, i.e.

$$X_t^i = \text{sign} \left[r_{t-12,t}^i \right]:$$

$$r_{t,t+1}^{TSMOM} = \frac{1}{N_t} \sum_{i=1}^{N_t} \text{sign} \left[r_{t-12,t}^i \right] \cdot \frac{\sigma_{tgt}}{\sigma_t^i} \cdot r_{t,t+1}^i \quad (9)$$

3.2. Turnover Dynamics

A long-only CVOL strategy involves frequent rebalancing due to the fact that the volatilities of the assets change from time to time and appropriate adjustment is necessary so that each asset maintains the same ex-ante target volatility. In contrast to this, a TSMOM strategy requires rebalancing because of two genuinely different effects: (i) because, similar to the CVOL strategy, the volatility of the portfolio constituents changes and (ii) because the trading rule of some assets changes from positive to negative and vice versa, due to the change in the direction of the trends.

Building on these observations, we next illustrate and disentangle the two channels through which portfolio turnover is affected: (i) the volatility channel and (ii) the trading rule channel (for TSMOM strategies only). In order to facilitate the exposition of the effects, we assume a single-asset paradigm and a single trading period defined by two rebalancing dates $t - 1$ and t .

First, consider a single-asset constant-volatility strategy, or, alternatively, a single-asset time-series momentum strategy, whose trading rule at dates $t - 1$ and t remains constant (either long or short). The turnover of the strategy would then be proportional to the change of the *reciprocal* of volatility. From Equation (4), we can deduce the marginal effect of volatility on portfolio turnover of a single-asset constant-volatility or time-series momentum strategy:

$$\text{turnover}_{vol}(t-1, t) \propto \left| \frac{1}{\sigma_t} - \frac{1}{\sigma_{t-1}} \right| = \left| \Delta \left(\frac{1}{\sigma_t} \right) \right| \quad (10)$$

Arguably, the smoother the transition between different states of volatility, the lower the turnover of a strategy. However, volatility is not directly observable, but instead it has to be estimated. The objective of the econometrician is to estimate σ_t at every rebalancing date. Realised volatility is estimated with error, that is $\hat{\sigma}_t = \sigma_t + \varepsilon_t$, where ε_t denotes the estimation error. Consequently, the turnover of the strategy is not only a function of the underlying volatility path, but more importantly of the error inherent in the estimation of the unobserved volatility path.

We hypothesise that larger in magnitude and time-varying estimation error results in over-trading and therefore in increased turnover in line with Fleming et al. (2003). Our related conjecture is that a more efficient volatility estimator can significantly reduce the turnover of a CVOL or TSMOM strategy and hence improve the performance of the strategies after accounting for transaction costs.

Apart from the volatility component, the rebalancing of a TSMOM strategy could alternatively be due to the switching of a position from long to short or vice versa. In order to focus on the marginal effect of the trading rule, assume that the volatility σ of an asset stays constant between the rebalancing dates $t - 1$ and t , but the position switches sign. The marginal effect of a trading rule on the turnover of a single-asset time-series momentum strategy is illustrated by the following relationship:

$$turnover_{rule}(t - 1, t) \propto \left| \frac{X_t}{\sigma} - \frac{X_{t-1}}{\sigma} \right| = \frac{|\Delta X_t|}{\sigma} \quad (11)$$

For a trading rule that takes only two values, such as the sign of the past return that only takes values $+1$ or -1 , $|\Delta X_t| = 2$, when the position switches sign from long to short or vice versa. In a more general setup that the trading rule has more than two states or even becomes a continuous function of past performance, the turnover of the TSMOM strategy would largely depend on the speed at which the trading rule changes states. The effect is also expected to be magnified for lower volatility assets, such as interest rate futures, since volatility appears in the denominator of Equation (11). This leads to the conjecture that a trading rule, which can avoid frequent swings between long and short positions and only instructs trading in the presence of strong price trends, can significantly reduce the turnover of a TSMOM strategy and therefore improve the performance of the strategy after accounting for transaction costs.

We empirically test the hypotheses relating to portfolio turnover reduction based on either the volatility estimator or the trading rule in Section 4.

3.3. Incorporating Pairwise Correlations

The construction of the TSMOM strategy in equations (8) and (8), which follows the standard specification used by other papers in the literature (Moskowitz et al. 2012, Hurst et al. 2013, Baltas and

Kosowski 2013) does not explicitly model the pairwise correlations between futures contracts as part of the weighting scheme. This potentially constitutes an important limitation for the strategy, especially in periods of increased asset co-movement, like the post-GFC period. One of the main methodological contributions of this paper is the extension in the formulation of the TSMOM strategy by taking into account the average pairwise correlation of portfolio constituents in an effort to improve the portfolio risk-return characteristics.

To do so we first investigate the interplay between the portfolio volatility and the pairwise correlations of portfolio constituents. Assume a portfolio of N assets with weights and volatilities denoted by w_i and σ_i for $i = 1, \dots, N$ respectively. To facilitate the notation, we drop the dependence on time in the following derivations. The portfolio volatility, σ_P , is trivially deduced as follows:

$$\sigma_P = \sqrt{\sum_{i=1}^N w_i^2 \sigma_i^2 + 2 \sum_{i=1}^N \sum_{j=i+1}^N w_i w_j \sigma_i \sigma_j \rho_{i,j}} \quad (12)$$

where $\rho_{i,j}$ denotes the pairwise correlation between assets i and j . The TSMOM strategy of equations (8) and (9) consists of assets whose weights are such that they all have ex-ante volatility equal to a pre-determined target level σ_{tgt} . In particular, each asset has an absolute weight/leverage factor equal to $\sigma_{tgt}/(N \cdot \sigma_i)$. Substituting the portfolio weights of Equation (12) with this quantity leads to the following result:

$$\begin{aligned} \sigma_P &= \sigma_{tgt} \sqrt{\sum_{i=1}^N \frac{1}{N^2} + 2 \sum_{i=1}^N \sum_{j=i+1}^N \frac{1}{N^2} \rho_{i,j}} \\ &= \frac{\sigma_{tgt}}{N} \sqrt{N + 2 \sum_{i=1}^N \sum_{j=i+1}^N \rho_{i,j}} \end{aligned} \quad (13)$$

The double summation $\sum_{i=1}^N \sum_{j=i+1}^N \rho_{i,j}$ is effectively the sum of all the elements of the upper right triangle of the correlation matrix of the assets. Normalising this quantity by the number of pairs formed by N assets (which can be trivially shown to be $\frac{N(N-1)}{2}$) results to the average pairwise correlation of the universe, $\bar{\rho}$:

$$\bar{\rho} = 2 \frac{\sum_{i=1}^N \sum_{j=i+1}^N \rho_{i,j}}{N(N-1)} \quad (14)$$

Solving for the double summation and substituting back into Equation (13) yields:

$$\sigma_P = \sigma_{tgt} \sqrt{\frac{1 + (N-1)\bar{\rho}}{N}} \quad (15)$$

The above result lies at the heart of diversification. Given that $\bar{\rho} \leq 1$, we deduce that $\sqrt{\frac{1+(N-1)\bar{\rho}}{N}} \leq 1$, and therefore that $\sigma_P \leq \sigma_{tgt}$. In other words, the fact that correlation across assets is empirically less than perfect results in a portfolio of assets with lower volatility than the target level of volatility of each asset. Thus, when correlation falls, diversification benefits increase and portfolio volatility drops further.

Following from Equation (15), we can introduce the average pairwise correlation as a factor that controls the target level of volatility of each asset. When average pairwise correlation increases (decreases) we would optimally want to lower (increase) the per asset target level of volatility. Solving Equation (15) for a dynamic level of target volatility for each asset results in:

$$\sigma_{tgt}(\bar{\rho}) = \sigma_P \sqrt{\frac{N}{1 + (N-1)\bar{\rho}}} \quad (16)$$

$$= \sigma_P \cdot CF(\bar{\rho}) \quad (17)$$

where

$$CF(\bar{\rho}) = \sqrt{\frac{N}{1 + (N-1)\bar{\rho}}} \quad (18)$$

denotes a correlation factor (CF) that adjusts the level of leverage applied to each portfolio constituent as a function of their average pairwise correlation.

Following the above, the generalised TSMOM strategy of Equation (8) can be accordingly adjusted by replacing the volatility target for each asset, σ_{tgt} with a time-varying target level of volatility that is determined by a target level of volatility for the *overall* strategy, $\sigma_{P,tgt}$ and a measure of the contemporaneous average pairwise correlation of the assets. This gives rise to the correlation-adjusted time-series momentum strategy (TSMOM-CF):

$$r_{t,t+1}^{TSMOM-CF} = \frac{1}{N_t} \sum_{i=1}^{N_t} X_t^i \cdot \frac{\sigma_{P,tgt}}{\sigma_t^i} \cdot CF(\bar{\rho}_t) \cdot r_{t,t+1}^i \quad (19)$$

We empirically study the effect of the correlation adjustment in Section 5, with a particular focus on

the post-GFC period, when pairwise correlations across assets and asset classes increased significantly, thus diminishing any diversification benefits.

4. Turnover Reduction

The purpose of this section is to empirically investigate the turnover implications of the two key determinants of a time-series momentum strategy, namely the volatility estimator and the trading rule. As motivated in the previous section, more efficient volatility estimators and more robust trading rules can be expected to reduce the turnover of the strategy and subsequently improve the performance after accounting for transaction costs.

4.1. The Effect of Volatility Estimator

Fleming et al. (2003) show that increasing the efficiency of volatility estimates can result in significant economic benefits for a risk-averse investor that dynamically rebalances a mean-variance optimised portfolio. The efficiency gain is achieved by switching from daily to high-frequency returns in order to estimate the conditional covariance matrix that is used in the optimisation. Extending this finding, we hypothesise that more efficient volatility estimates can significantly reduce portfolio turnover and consequently improve the net of transaction costs profitability of CVOL and TSMOM strategies.

The ordinary measure of volatility is the standard deviation of past daily close-to-close returns (STDEV, hereafter), which, even though an unbiased estimator, it only makes use of daily closing prices and therefore is subject to large estimation error when compared to volatility estimators that make use of intraday information. In the absence of high-frequency data in our dataset, we attempt to improve the estimation efficiency of a close-to-close daily volatility estimator by using intraday open, high, low and close daily prices. The volatility estimators that make use of open, high, low and close prices are known in the literature as range⁶ estimators and have been shown to offer additional robustness against mi-

⁶The term “range” refers to the daily high-low price difference and its major advantage is that it can successfully capture the high volatility of an erratically moving price path intra-daily, which happens to exhibit similar opening and closing prices and therefore a low daily return. As an indicative example, on Tuesday, August 9, 2011, most major exchanges demonstrated price erratic behaviour, following the previous day’s large losses and the downgrade of the US’s sovereign debt rating from AAA to AA+ by Standard & Poor’s late on Friday, August 6, 2011. On that Tuesday, FTSE100 exhibited intra-daily a 5.48% loss and a 2.10% gain compared to its opening price, before closing 1.89% up. An article in the Financial Times entitled “Investors

crostructure noise such as bid-ask bounce and asynchronous trading and therefore increase the efficiency of the estimation (Alizadeh, Brandt and Diebold 2002).

A multitude of range estimators have been suggested in the literature by Parkinson (1980), Garman and Klass (1980), Rogers and Satchell (1991), and Yang and Zhang (2000), which have been empirically shown (see for example Brandt and Kinlay 2005, Shu and Zhang 2006) to reduce the estimation error of a conventional daily volatility estimator, like the standard deviation of past returns. Out of these estimators, the Yang and Zhang (2000) estimator (YZ, hereafter) is the most efficient and the only to be independent of both the overnight jump (i.e. the price change between the previous day's close and the next day's opening price) and the drift of the price process. For that reason and for the purposes of our analysis we focus solely on the added benefit of more efficient volatility estimates, as these are offered by the the YZ estimator.⁷

The YZ estimator is defined as a linear combination of three volatility estimators: the standard deviation of past close-to-close daily logarithmic returns (i.e. the conventional STDEV estimator), the standard deviation of past overnight (close-to-open) logarithmic returns and the Rogers and Satchell (1991) (RS, hereafter) range estimator⁸. In particular, the YZ volatility of an asset at the end of month t (assuming some estimation period) is given by:

$$\sigma_{YZ}^2(t) = \sigma_{OJ}^2(t) + k \cdot \sigma_{STDEV}^2(t) + (1 - k) \cdot \sigma_{RS}^2(t) \quad (20)$$

where $\sigma_{OJ}(t)$ denotes the overnight jump estimator. The parameter k is chosen so that the variance of the estimator is minimised and is shown by Yang and Zhang (2000) to be a function of the number of days

shaken after rollercoaster ride" on August 12 mentions that "...the high volatility in asset prices has been striking. On Tuesday, for example, the FTSE100 crossed the zero per cent line between being up or down on that day at least 13 times...".

⁷In undocumented results, we have additionally evaluated the performance of the less efficient range estimators by Parkinson (1980), Garman and Klass (1980) and Rogers and Satchell (1991). These results are available upon request.

⁸Rogers and Satchell (1991) are the first to introduce an unbiased estimator that allows for a non-zero drift in the price process, but their estimator does not account for the overnight jump (see also Rogers, Satchell and Yoon 1994). Their estimator is 6.2 times more efficient than STDEV. The RS volatility of an asset over the course of a single day τ is given by:

$$\sigma_{RS}^2(\tau) = h(\tau) [h(\tau) - c(\tau)] + l(\tau) [l(\tau) - c(\tau)]$$

where $h(\tau)$, $l(\tau)$ and $c(\tau)$ denote the logarithmic difference between the high, low and closing prices respectively and the opening price. The RS volatility of an asset at the end of month t , assuming a certain estimation period is equal to the average daily RS volatility over this period.

used in the estimation.⁹ The YZ estimator is $1 + \frac{1}{k}$ times more efficient than STDEV; this expression is maximised for a 2-day estimator, when YZ is almost 14 times more efficient than STDEV. For our purposes, a monthly YZ estimator with -on average- 21 daily returns would be 8.2 times more efficient than the monthly STDEV estimator.

4.1.1. Performance Evaluation

We start our analysis by exploring the effects of a more efficient volatility estimator on the turnover of CVOL and TSMOM portfolios that are constructed as in equations (6) and (9) respectively. In line with Moskowitz et al. (2012) and Baltas and Kosowski (2013) we use a volatility target for each asset equal to $\sigma_{tgt} = 40\%$. This choice is motivated in these studies by the fact that it generates ex-post TSMOM portfolio volatilities that are comparable to those of commonly used factors such as those constructed by Fama and French (1993) and Asness et al. (2013).

Table II presents out-of-sample performance statistics for CVOL (Panel A) and TSMOM (Panel B) strategies that employ a different volatility estimator at a time. The period of volatility estimation is fixed at one month; robustness results using longer windows of estimation follow later in this section. The last column of the table reports statistics for a hypothetical strategy that uses the ex-post realised volatility over the holding month to ex-ante scale the futures positions. This strategy cannot be implemented in real-time and only constitutes a benchmark for the purpose of our analysis; for this reason, we label it as the “perfect forecast” strategy (PF, hereafter).

[Table II about here]

As expected, the two different volatility estimators, STDEV and YZ, do not have an economically significantly different effect on the performance of the CVOL or TSMOM strategies, at least before accounting for transaction costs. The risk-adjusted returns are around 0.60 for the long-only strategies and around 0.82 for the time-series momentum strategies. However, our focus is on the effect of the more

⁹The parameter k is chosen using the following equation:

$$k = \frac{0.34}{1.34 + \frac{N_D + 1}{N_D - 1}}$$

where N_D denotes the number of days in the estimation period.

efficient volatility estimator on portfolio turnover, given that the level of estimation noise is expected to be relatively lower. We find that the more efficient YZ estimator reduces the turnover by almost one tenth; 13% for the CVOL strategy and 8% for the TSMOM strategy. This result is indeed in line with our conjecture; the turnover benefit should translate in practice into lower trading costs and therefore larger net returns for the strategy.

Comparing the results of the implementable strategies to the PF benchmark, it is obvious that the strategy based on the perfect forecast delivers larger risk-adjusted performance with a before-transaction costs Sharpe ratio of 0.89 (CVOL) and 1.30 (TSMOM), which are significantly different from the Sharpe ratios of the rest of the strategies as deduced by the very low p-values of the Ledoit and Wolf (2008) statistical test. The rejection of the null of equality in Sharpe ratios shows that there is room for improvement in terms of accurately forecasting increases (decreases) in volatility and therefore better timing the downscaling (upscaling) of positions before an impending drawdown (uptrend). However, this task is beyond the scope of this paper. Our main objective is to show that increased estimation efficiency can significantly reduce the turnover and therefore the transaction costs of a CVOL or TSMOM strategy and not to forecast future realised volatility.

Given the documented turnover benefit of the more efficient YZ volatility estimator, next we investigate whether the turnover reduction is pervasive across all portfolio constituents or whether instead the result is dominated by a few assets. For that reason, we use monthly STDEV and YZ volatility estimates for all 75 future contracts of our dataset and calculate the time-series average absolute first order difference in the reciprocal of volatility estimates, which is a quantity that, as shown in Equation (10), directly affects the turnover of a strategy. For that purpose, we call this statistic the “Volatility Turnover”:

$$\text{Volatility Turnover}(i, \text{estimator}) = \frac{1}{\#months} \sum_m \left| \frac{1}{\sigma_{i,\text{estimator}}(t_m, t_{m+1})} - \frac{1}{\sigma_{i,\text{estimator}}(t_{m-1}, t_m)} \right| \quad (21)$$

where estimator = {STDEV, YZ}. In principle, a more efficient volatility estimator should reduce the volatility turnover statistic for each asset. Figure 1 presents the percentage drop in the volatility turnover statistic when switching from the STDEV estimator to the YZ estimator. In other words, we plot the value $100 \cdot \left(\frac{\text{Volatility Turnover}(i, \text{YZ})}{\text{Volatility Turnover}(i, \text{STDEV})} - 1 \right)$ for each asset i .

[Figure 1 about here]

The empirical evidence is very strong. Across all 75 contracts, without any exception, the time-series average change in the reciprocal of volatility is reduced when the more efficient YZ volatility estimator is used. The effects are more pronounced for low volatility contracts, like the interest rate contracts, but even for equity contracts the average drop is above 10%, with the maximum drop being exhibited for the S&P500 contract at about 26%. These results suggest that the large error variance of the STDEV estimator is the main reason for excessive overtrading in a CVOL or TSMOM strategy.

4.1.2. Robustness Test - The Effect of Estimation Period

In Table II we quantified the economic benefit of a more efficient volatility estimator in terms of turnover reduction using an estimation window of one month. Next, we examine whether the choice of the volatility estimation window affects the marginal benefit of using the YZ range estimator over the standard STDEV estimator. Figures 2 (for CVOL strategies) and 3 (for TSMOM strategies) report different performance statistics, including the turnover benefit, for different sizes of the estimation window ranging from one to twelve months.

[Figure 2 about here]

[Figure 3 about here]

Figure 3 shows that the Sharpe ratio of the TSMOM strategy, before transaction costs, is maximised when using a three-month estimation window. Although this recommendation is based on empirical results, it lends support to the choice of a three-month volatility estimation window in Baltas and Kosowski (2013) and a 60-business day centre of mass of the weights of the exponentially weighted volatility estimator in Moskowitz et al. (2012) and Hurst et al. (2013). Based on these results, we will use an estimation window of three months in the remainder of the paper.

4.2. The Effect of Trading Rule

The second part of the turnover reduction analysis relates to the trading rule that is employed by the TSMOM strategy. In this section, we explore the mechanics of two different trading rules and investigate

how they affect the turnover and the performance of the strategy. In particular, we compare the standard time-series momentum rule (sign of the past return) with a rule that identifies statistically significant trends and therefore allows exiting a position in the absence of one.

Return Sign (SIGN): The ordinary measure of past performance that has been used in the literature (Moskowitz et al. 2012, Hurst et al. 2013, Baltas and Kosowski 2013) as well as in our paper so far is the sign of the past 12-month return. A positive (negative) past return dictates a long (short) position:

$$\text{SIGN}_t^{12M} = \text{sign}[r_{t-12,t}] = \begin{cases} +1, & r_{t-12,t} \geq 0 \\ -1, & \text{otherwise} \end{cases} \quad (22)$$

Time-Trend t-statistic (TREND): Another way to capture the trend of a price series is by fitting a time-trend on the past 12-month daily futures log-price series using least-squares. The momentum trading rule can then be determined by the significance of the slope coefficient of the trend fit. Assume the linear regression model:

$$\log(F_\tau) = \alpha + \beta \cdot \tau + e_\tau, \quad \tau \in [t_{t-12}, t] \quad (23)$$

The significance of the time-trend is determined by the Newey and West (1987) t-statistic of β , $t(\beta)$, and the cutoff points for the long/short position of the trading rule are chosen to be +2/-2 respectively:

$$\text{TREND}_t^{12M} = \begin{cases} +1, & \text{if } t(\beta) > +2 \\ -1, & \text{if } t(\beta) < -2 \\ 0, & \text{otherwise} \end{cases} \quad (24)$$

The TREND rule effectively instructs staying out of certain assets if no statistically significant price trends are identified. Our aim is to investigate how the sparse trading rule affects the performance and turnover of the TSMOM strategy.¹⁰ This, in turn, allows us to address the question of whether statistically significant price trends are the main performance drivers of the TSMOM strategy.

¹⁰Clearly, more sophisticated methodologies in constructing momentum strategies can be devised, but our objective is to maintain a simple and tractable framework and avoid data mining.

4.2.1. Return Predictability

Following Moskowitz et al. (2012) and Baltas and Kosowski (2013), we first assess the amount of in-sample return predictability that is inherent in lagged excess returns or lagged trading rules by running the following pooled panel regressions:

$$\frac{r_{t-1,t}}{\sigma_{t-1}} = \alpha + \beta_\lambda \cdot \frac{r_{t-\lambda-1,t-\lambda}}{\sigma_{t-\lambda-1}} + \varepsilon_t \quad (25)$$

and

$$\frac{r_{t-1,t}}{\sigma_{t-1}} = \alpha + \beta_\lambda \cdot X_{t-\lambda}^{1M} + \varepsilon_t \quad (26)$$

where λ denotes the lag that ranges between 1 and 60 months and the lagged one month rule $X_{t-\lambda}^{1M}$ is either the $\text{SIGN}_{t-\lambda}^{1M}$ or the $\text{TREND}_{t-\lambda}^{1M}$ rule.

The regressions (25) and (26) are estimated for each lag by pooling together all T_i (where $i = 1, \dots, N$) monthly returns/trading rules for the $N = 75$ contracts. We are interested in the t -statistic of the coefficient β_λ for each lag. Large and significant t -statistics support the hypothesis of time-series return predictability. The t -statistics $t(\beta_\lambda)$ are computed using standard errors that are clustered by time and asset,¹¹ in order to account for potential cross-sectional dependence (correlation between contemporaneous returns of the contracts) or time-series dependence (serial correlation in the return series of each individual contract). Briefly, the variance-covariance matrix of the regressions (25) and (26) is given by (Cameron, Gelbach and Miller 2011, Thompson 2011):

$$V_{\text{TIME\&ASSET}} = V_{\text{TIME}} + V_{\text{ASSET}} - V_{\text{WHITE}}, \quad (27)$$

where V_{TIME} and V_{ASSET} are the variance-covariance matrices of one-way clustering across time and asset respectively, and V_{WHITE} is the White (1980) heteroscedasticity-robust OLS variance-covariance matrix. In fact, Petersen (2009) shows that when $T \gg N$ ($N \gg T$) then standard errors computed via one-way clustering by time (by asset) are close to the two-way clustered standard errors; nevertheless, one-way clustering across the “wrong” dimension produces downward biased standard errors, hence inflating

¹¹Petersen (2009) and Gow, Ormazabal and Taylor (2010) study a series of empirical applications with panel datasets and recognise the importance of correcting for both forms of dependence.

the resulting t -statistics and leading to over-rejection rates of the null hypothesis. Our panel dataset is unbalanced as not all assets have the same number of monthly observations. On average, we have $\bar{T} = \frac{1}{N} \sum_{i=1}^N T_i \cong 319$ months of data per asset. We can therefore argue that $\bar{T} > N$ and we document in unreported results (available upon request) that two-way clustering or one-way clustering by time (i.e. estimating T cross-sectional regressions as in Fama and MacBeth 1973) produces similar results, whereas clustering by asset produces inflated t -statistics that are similar to OLS t -statistics. Two-way clustering is also used by Baltas and Kosowski (2013), who study the return predictability over monthly, weekly and daily frequencies, whereas one-way clustering by time is used by Moskowitz et al. (2012).

Figure 4 presents the two-way clustered t -statistics $t(\beta_\lambda)$ for regressions (25) and (26) and lags $\lambda = 1, 2, \dots, 60$ months. The t -statistics are almost always positive for the first twelve months for all regressor choices, hence indicating strong momentum patterns of past year's returns. Moreover, the fact that the TREND rule is sparsely active does not seem to affect its return predictability, which also remains statistically strong for the first twelve months. Apparently, it is the statistical significance of the price trends that drives the documented momentum behaviour. In line with Moskowitz et al. (2012) and Baltas and Kosowski (2013), there exist statistically strong signs of return reversals after the first year¹² that subsequently attenuate and only seem to gain some significance for a lag of around three years.

[Figure 4 about here]

4.2.2. Performance Evaluation

Similar to the analysis in Table II, which studies the impact of the volatility estimator choice on portfolio turnover, in Table III we examine the economic value of using the SIGN or TREND trading rule in a TSMOM strategy. The results show that the SIGN trading rule has a slightly higher Sharpe ratio than the TREND trading rule before transaction costs. However, the Ledoit and Wolf (2008) p -value shows that the Sharpe ratios of 1.04 and 0.99 respectively are not statistically different from each other. Our focus is on the turnover reduction that can potentially be achieved by a more accurate indicator of price trends.

¹²Part of this severe transition from largely positive and significant t -statistic to largely negative and significant t -statistic after the lag of twelve months can be potentially attributed to seasonal patterns in the commodity futures returns. In undocumented results, we repeat the pooled panel regression only on commodity contracts, after removing contracts that for various reasons might exhibit seasonality, like the agricultural and energy contracts. In general the patterns become relatively less pronounced, but our conclusions remain qualitatively the same and the momentum/reversal transition is still apparent.

Indeed, the turnover generated by the SIGN rule is almost three times as high as that of the TREND trading rule. This is a very strong result, because it implies that the TREND rule leads to a similar before transaction costs Sharpe ratio, but only requires one third of the trading and associated cost. To a large extent, this is the consequence of staying out of certain assets in periods of statistically insignificant trends.

[Table III about here]

Having studied the aggregate strategy, we next study the effect on the before costs Sharpe ratio and turnover from switching between the two trading rules at the asset level. Panel A of Figure 5 presents the Sharpe ratio of single-asset time-series momentum strategies that use the SIGN rule, as well as the change in the Sharpe ratio (i.e. ΔSharpe) from using the TREND instead. On average ΔSharpe is insignificant as the TREND rule leads to an increase for some contracts and a decrease for others. The reductions appear to be concentrated among fixed income and commodities contracts. Panel B of Figure 5 shows the effect on turnover from switching between SIGN and TREND rules and supports earlier conclusions that using the latter has an economically large effect on performance net of transaction costs. The reduction in turnover is pervasive across all assets and is numerically around two thirds for most of them (ranging overall between 55% and 85%). This shows that the turnover benefit for the aggregate strategy is not due to a small number of assets.

[Figure 5 about here]

5. The Recent Underperformance of Time-series Momentum Strategies and the Effect of Pairwise Correlations

The purpose of this section is to investigate the dependence of the performance of the TSMOM strategy on the level of the pairwise correlations of portfolio constituents and shed light on the poor performance of the strategy after 2008. Our analysis is motivated by the results of Baltas and Kosowski (2013), who, after finding no significant evidence of capacity constraints in the performance of the strategy, argue that the underperformance can potentially be attributed to the lack of significant price trends or an

increased level of correlation across assets of different asset classes in the period from 2008 to 2013.¹³ We first empirically document evidence supporting these two claims, namely the lack of significant price trends over the most recent period and the increase in the correlations. Subsequently, we investigate the economic benefit from incorporating pairwise correlations in the weighting scheme of the TSMOM strategy.

5.1. Price Trend Significance and Pairwise Correlations

Panel A of Figure 6 shows the (12-month moving average) percentage of contracts for which the SIGN and TREND trading rules have the same value (either 1 or -1) at the end of each month. We notice a drop of the statistic towards the end of the sample period, which implies that the TREND rule is likely to return more 0's (hence documenting fewer significantly trending assets). In fact, Panel B of Figure 6 presents the (12-month moving average) percentage of assets with TREND=0 at the end of each month, i.e. the proportion of the investable universe that shows no signs of significant price trend. We find that after 2008 the number of contracts without a significant price trend increases significantly and almost doubles; the statistic goes from around 9% in the beginning of 2009 up to around 18% in the beginning of 2013. This absence of strong momentum patterns is, therefore, one possible reason for the recent performance drop of the TSMOM strategy.

[Figure 6 about here]

In line with our finding of the lack of significant price trends in the post-GFC period, the average level of pairwise correlation has dramatically increased in 2009 and remained at higher levels than the historical average since then. This is evident from Panel A of Figure 7 which shows the average pairwise correlation across all contracts in our sample using a six-month rolling estimation window (the values prior to 1983 should be treated with caution as during that period the dataset consists of less than 20 traded contracts). In an environment of increased correlations, the portfolio construction methodology that is typically employed in TSMOM strategies can be suboptimal, as it lacks an adjustment for the aggregate level of co-movement.

¹³See for example, "CTA trend-followers suffer in market dominated by intervention" by Emma Cusworth, Hedge Funds Review, 10 October 2013.

An important research question is whether incorporating information from the correlation matrix of the assets into the portfolio construction can render the strategy more robust in periods of increased co-movement. As explained in Section 3, this can be achieved by using the contemporaneous level of average pairwise correlation to dynamically adjust the target level of volatility of each asset in the TSMOM strategy. The correlation-adjusted TSMOM strategy of Equation (19) is repeated below for convenience, using the SIGN trading rule:

$$r_{t,t+1}^{TSMOM-CF} = \frac{1}{N_t} \sum_{i=1}^{N_t} \text{sign}[r_{t-12,t}^i] \cdot \frac{\sigma_{P,tgt}}{\sigma_i} \cdot CF(\bar{\rho}_t) \cdot r_{t,t+1}^i \quad (28)$$

The correlation factor (CF) depends on the average pairwise correlation (see Equation (18)) and is plotted in Panel B of Figure 7. As expected, incorporating the correlation factor would call for a reduction in the leverage that is employed during the period 2009-2013 when pairwise correlations increase and diversification benefits diminish. We next evaluate the performance of the correlation-adjusted TSMOM strategy. It is of particular interest whether any benefit is more pronounced after the recent financial crisis since correlations have been relatively high during that time period compared to their historical average.

[Figure 7 about here]

5.2. Performance Evaluation

Table IV reports performance statistics for the standard TSMOM strategy with a per asset target level of volatility $\sigma_{tgt} = 40\%$ and the correlation-adjusted strategy with a portfolio target level of volatility $\sigma_{P,tgt} = 12\%$. These choices of asset and portfolio target volatility are inconsequential and are justified by the arguments of Moskowitz et al. (2012), who similarly choose the asset target level of volatility to be 40%, so that the overall TSMOM strategy exhibits an ex-post portfolio volatility of 12% for their respective sample period (1985-2009), which, in turn, matches roughly the level of volatility of the Fama and French (1993) and Asness et al. (2013) factors.

[Table IV about here]

Several interesting insights emerge from the analysis. Over the full sample, the correlation adjustment does not lead to a significant increase in the Sharpe ratio which increases marginally from 1.04 to 1.05. However, the incorporation of correlation into the weighting scheme renders the portfolio more robust to crash/downside risk. This is evidenced by the relatively larger values of performance ratios that measure risk using downside volatility (the so-called Sortino ratio; see Sortino and Van Der Meer 1991) or by the maximum drawdown (the so-called Calmar ratio; see Young 1991). Taking into account pairwise correlations can positively affect the diversification benefits of the portfolio. However, this benefit does not come without cost. As expected the effect of modelling and incorporating correlation does increase portfolio turnover (from 28.89% to 41.23% across the entire sample period).

Note, however, that incorporating the information from the correlation matrix has a relatively stronger effect during the period after the 2008 financial crisis (January 2009 to February 2013); see Panel B of Table IV. This period has been characterised by elevated levels of co-movement across asset classes and therefore by diminished diversification benefits. During this period, incorporating correlations in the construction of a TSMOM strategy has an economically significant effect since it more than doubles the before-transaction cost Sharpe ratio from 0.14 to 0.29. Other risk-adjusted performance ratios also increase significantly. Figure 8 summarises the benefit from incorporating the correlation adjustment in the time-series momentum strategy both across the entire sample period and across the post-crisis period.

[Figure 8 about here]

Overall, our results highlight the role that the increased pairwise correlations have played in the recent performance of simple unadjusted TSMOM strategies. The implication for fund managers and investors is two-fold. First, allowing correlation to determine portfolio weights and dynamically adjust the level of leverage that is employed can be beneficial during periods of high correlations. Second, adjusting for correlation can increase trading costs. Whether the risk-adjusted performance benefits survive after taking into account the higher turnover and respective transaction costs remains to be tested. The next section of the paper provides a set of after-costs results using a simple cost approximation model.

6. Trading Costs Implications

Our results so far have shown that different volatility estimators, trading rules and correlation adjustments can substantially impact portfolio turnover, which, as a result, can potentially affect the performance of a strategy after accounting for transaction costs. In this section, we present a simple model for the approximation of transaction costs and evaluate the after-costs performance of the various variants of the time-series momentum strategy.

Transaction costs have several components and these are typically classified into implicit and explicit transaction costs (Harris 2002). Explicit transaction costs include brokerage commissions, market fees, clearing and settlement costs, and taxes/stamp duties. Implicit transaction costs refer to costs that are not explicitly included in the trade price and therefore have to be estimated. They mainly consist of bid-ask spreads, market impact, operational opportunity costs, market timing opportunity costs and missed trade opportunity costs.

Market participants estimate implicit transaction costs using specified price benchmark methods and econometric transaction cost estimation methods (Harris 2002). These costs depend primarily on the characteristics of any given trade relative to prevailing market conditions and include factors such as the order size as well as the asset's daily trading volume and volatility. Several papers focus on estimating transaction costs in cash equities (see for example Jones and Lipson 1999) using institutional equity order data. We do not have access to such data for the futures contracts that we study in this paper. Moreover, even with such data available, it would be beyond the scope of our paper to estimate implicit and explicit transaction costs in detail.

For the purpose of our analysis, we therefore provide a simple approximation for transaction costs based on the turnover of the strategy. In this respect, we follow Frazzini et al. (2012) and estimate the realised costs of each strategy as a product of the turnover and an assumed level of market impact:

$$\text{Realised Costs} = \text{Turnover} * \text{Market Impact} \quad (29)$$

Frazzini et al. (2012) use proprietary data of a large institution on portfolio holdings and execution prices for the construction and rebalancing of various equity portfolios (e.g. momentum, value) and

estimate average market impact costs of around 20 basis points. They argue that these cost estimates are relatively lower to what earlier academic research has found and this is largely driven by the fact that they have access to this proprietary dataset of actual execution prices of a large institutional investor instead of, say, the volume-weighted average price (VWAP) that is typically used to approximate the execution price for the average investor.

For the purpose of our analysis, we assume a constant level of market impact and therefore the realised costs become proportional to portfolio turnover. We acknowledge that this simplification constitutes an important limitation of our approach, as in practice the market impact depends, among other things, on the order size with the relationship being typically convex. Time-series momentum strategies are typically constructed using futures contracts, which are generally more liquid and exhibit lower transaction costs than stocks. Using the findings of Frazzini et al. (2012) as guideline, we estimate realised costs for the time-series momentum strategies using two different levels for the average market impact: a realistic scenario of 10 basis points and a conservative scenario of 50 basis points. The annualised realised costs as well as the after-costs Sharpe ratio of the different strategies that we study in this paper are reported in Table V. We note that the findings in this section are intended to be illustrative since they depend on the assumptions related to the cost approximation in Equation (29).

[Table V about here]

Panels A and B of Table V show that more efficient volatility estimation and price trend detection can significantly reduce the respective turnover and therefore the rebalancing costs of the strategies. Given the cost approximation of Equation (29), the associated trading costs are proportional to portfolio turnover and therefore the relative percentage drop in the realised costs from using either a more efficient volatility estimator or the TREND rule is identical to the respective turnover drop that is reported in Tables II and III. The cost reduction from using a more efficient volatility estimator such as the YZ estimator is around 10-20% and that from using a sparse trading signal is around 66%. Nevertheless, the impact on the risk-adjusted returns of the strategy is not as large, as could have reasonably been expected given the magnitude of the change in turnover.

This finding appears surprising since portfolio turnover is a critical parameter in a dynamic portfolio rebalancing framework and any significant reduction of it can be expected to benefit the net of costs

performance of the strategy. However, even though the theoretically superior volatility estimation and price trend detection reduce the turnover of the strategy, it is not obvious that one can obtain significant economic benefits from using them. One potential explanation for this finding could be the simplicity of the trading costs model that we use, which might fail to capture the dynamics of rebalancing a multi-asset portfolio such as time-series momentum strategy. Future research that optimally employs detailed data on actual futures-related transaction costs could be used to refine the analysis further.

Focusing on the effect of correlation in Panel C of Table V, we find that the estimated transaction costs increase by around 40-50% when the correlation adjustment is employed. However, the improvement in the risk-adjusted performance of the correlation-adjusted strategy survives the larger costs, especially in the post-crisis period. If transaction costs are ignored, the Sharpe ratio of the unadjusted time-series momentum strategy during the post-GFC period is 0.14, while that of the correlation-adjusted strategy is 0.29. If we adjust for transaction costs assuming costs of 10 (50) basis points, then the Sharpe ratio of the strategies falls to 0.12 and 0.26 (0.05 and 0.14) respectively.

Overall, our analysis shows that adjusting the leverage of the portfolio based on the level of the pairwise correlations of its constituents can have an economically significant impact on the risk-adjusted performance both before and, most importantly, after the incorporation of realistic transaction costs. In other words, the rebalancing costs do not completely swamp the improved diversification property of the portfolio.

7. Concluding Remarks

The objective of this paper is twofold. First, we investigate the effect of the efficiency of volatility estimation and of the choice of momentum trading rule on portfolio turnover and the performance of time-series momentum strategies before and after transaction costs. Second, we study whether incorporating correlations can further help improve the performance of the strategy and shed light on the source of the poor performance of time-series momentum strategies in the years following the GFC.

We find that more efficient volatility estimates, like those provided by a range-based estimator, as well as trading rules that instruct trading only in the presence of statistically significant price trends can

substantially reduce the turnover of the time-series momentum strategy and resulting rebalancing costs. The benefit is pervasive across the entire list of assets that are typically included in such strategies.

In an effort to shed light on the recent underperformance of time-series momentum strategies in the period following the GFC, we introduce a correlation factor in the weighting scheme and extend the standard time-series momentum strategy. By dynamically adjusting the leverage that is employed as a function of the average pairwise correlation, this adjustment is shown to improve the diversification benefits and the performance of the strategy by safeguarding against crash risk. The performance benefit is significant and pronounced over the most recent post-crisis period, 2009-2013, during which pairwise correlations across assets and asset classes dramatically increased and diminished diversification benefits. The improvement in the performance remains robust to the incorporation of transaction costs.

Overall, our results shed light on the drivers of the recent underperformance of CTA funds and indicate ways to improve the performance of time-series momentum strategies either by estimating volatility and price trends more efficiently or by improving the diversification properties of the portfolio.

References

- Alizadeh, S., Brandt, M. W. and Diebold, F. X.: 2002, Range-based estimation of stochastic volatility models, *Journal of Finance* **57**(3), 1047–1091.
- Asness, C. S., Moskowitz, T. J. and Pedersen, L. H.: 2013, Value and momentum everywhere, *Journal of Finance* **68**(3), 929–985.
- Baltas, A. N. and Kosowski, R.: 2013, Momentum strategies in futures markets and trend-following funds, *SSRN eLibrary*.
- Barroso, P. and Santa-Clara, P.: 2015, Momentum has its moments, *Journal of Financial Economics* **116**(1), 111–120.
- Bessembinder, H.: 1992, Systematic risk, hedging pressure, and risk premiums in futures markets, *Review of Financial Studies* **5**(4), 637–667.
- Bessembinder, H.: 1993, An empirical analysis of risk premia in futures markets, *Journal of Futures Markets* **13**(6), 611–630.
- Brandt, M. W. and Kinlay, J.: 2005, Estimating historical volatility, *Research Article, Investment Analytics*.

- Cameron, A. C., Gelbach, J. B. and Miller, D. L.: 2011, Robust inference with multiway clustering, *Journal of Business and Economic Statistics* **29**(2), 238–249.
- Daniel, K. and Moskowitz, T.: 2015, Momentum crashes, *Columbia Business School Research Paper* .
- de Roon, F. A., Nijman, T. E. and Veld, C.: 2000, Hedging pressure effects in futures markets, *Journal of Finance* **55**(3), 1437–1456.
- Fama, E. F. and French, K. R.: 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* **33**(1), 3–56.
- Fama, E. F. and MacBeth, J.: 1973, Risk and return: Some empirical tests, *Journal of Political Economy* **81**(3), 607–636.
- Fleming, J., Kirby, C. and Ostdiek, B.: 2001, The economic value of volatility timing, *Journal of Finance* **56**(1), 329–352.
- Fleming, J., Kirby, C. and Ostdiek, B.: 2003, The economic value of volatility timing using realized volatility, *Journal of Financial Economics* **67**(3), 473–509.
- Frazzini, A., Israel, R. and Moskowitz, T. J.: 2012, Trading costs of asset pricing anomalies, *SSRN eLibrary* pp. 14–05.
- Fuertes, A., Miffre, J. and Rallis, G.: 2010, Tactical allocation in commodity futures markets: Combining momentum and term structure signals, *Journal of Banking and Finance* **34**(10), 2530–2548.
- Garman, M. B. and Klass, M. J.: 1980, On the estimation of security price volatilities from historical data, *Journal of Business* **53**(1), 67–78.
- Gorton, G. B., Hayashi, F. and Rouwenhorst, K. G.: 2007, The fundamentals of commodity futures returns, *NBER Working Paper* .
- Gorton, G. and Rouwenhorst, K. G.: 2006, Facts and fantasies about commodities futures, *Financial Analysts Journal* **62**(2), 47–68.
- Gow, I. D., Ormazabal, G. and Taylor, D. J.: 2010, Correcting for cross-sectional and time-series dependence in accounting research, *The Accounting Review* **85**(2), 483–512.
- Hallerbach, W. G.: 2012, A proof of the optimality of volatility weighting over time, *Journal of Investment Strategies* **1**(4), 87–99.
- Harris, L.: 2002, *Trading and Exchanges: Market Microstructure For Practitioners*, Oxford University Press.
- Hurst, B., Ooi, Y. H. and Pedersen, L. H.: 2013, Demystifying managed futures, *Journal of Investment Management* **11**(3), 42–58.

- Ilmanen, A. and Kizer, J.: 2012, The death of diversification has been greatly exaggerated, *Journal of Portfolio Management* **38**(3), 15–27.
- Jegadeesh, N. and Titman, S.: 1993, Returns to buying winners and selling losers: Implications for stock market efficiency., *Journal of Finance* **48**(1), 65–91.
- Jegadeesh, N. and Titman, S.: 2001, Profitability of momentum strategies: An evaluation of alternative explanations, *Journal of Finance* **56**(2), 699–720.
- Jones, C. and Lipson, M.: 1999, Execution costs of institutional equity orders, *Journal of Financial Intermediation* **8**(3), 123–140.
- Kirby, C. and Ostdiek, B.: 2012, Its all in the timing: simple active portfolio strategies that outperform naive diversification, *Journal of Financial and Quantitative Analysis* **47**(02), 437–467.
- Ledoit, O. and Wolf, M.: 2008, Robust performance hypothesis testing with the sharpe ratio, *Journal of Empirical Finance* **15**(5), 850–859.
- Miffre, J. and Rallis, G.: 2007, Momentum strategies in commodity futures markets, *Journal of Banking and Finance* **31**(6), 1863–1886.
- Moskowitz, T., Ooi, Y. H. and Pedersen, L. H.: 2012, Time series momentum, *Journal of Financial Economics* **104**(2), 228 – 250.
- Newey, W. K. and West, K. D.: 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* **55**(3), 703–708.
- Parkinson, M.: 1980, The extreme value method for estimating the variance of the rate of return, *Journal of Business* **53**(1), 61–65.
- Pesaran, M., Schleicher, C. and Zaffaroni, P.: 2009, Model averaging in risk management with an application to futures markets, *Journal of Empirical Finance* **16**(2), 280–305.
- Petersen, M. A.: 2009, Estimating standard errors in finance panel data sets: Comparing approaches, *Review of Financial Studies* **22**(1), 435–480.
- Pirrong, C.: 2005, Momentum in futures markets, *SSRN eLibrary* .
- Rogers, L. C. G. and Satchell, S. E.: 1991, Estimating variance from high, low and closing prices, *Annals of Applied Probability* **1**(4), 504–512.
- Rogers, L. C. G., Satchell, S. E. and Yoon, Y.: 1994, Estimating the volatility of stock prices: a comparison of methods that use high and low prices, *Applied Financial Economics* **4**(3), 241–247.
- Shu, J. and Zhang, J. E.: 2006, Testing range estimators of historical volatility, *Journal of Futures Markets* **26**(3), 297–313.

- Sortino, F. A. and Van Der Meer, R.: 1991, Downside risk, *Journal of Portfolio Management* **17**(4), 27–31.
- Thompson, S. B.: 2011, Simple formulas for standard errors that cluster by both firm and time, *Journal of Financial Economics* **99**(1), 1–10.
- White, H.: 1980, A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity, *Econometrica* **4**, 817–838.
- Yang, D. and Zhang, Q.: 2000, Drift-independent volatility estimation based on high, low, open, and close prices, *Journal of Business* **73**(3), 477–491.
- Young, T.: 1991, Calmar ratio: A smoother tool, *Futures* **20**(1), 40.

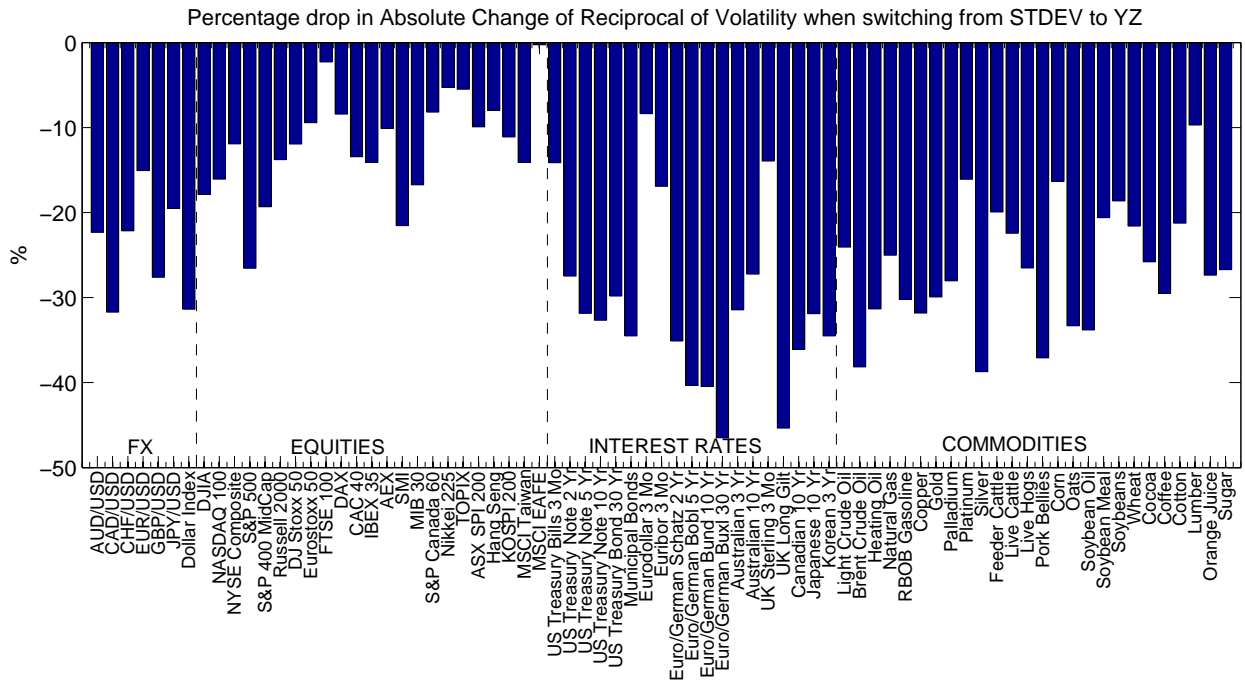


Figure 1: *Effect of Volatility Estimator choice on Reciprocal of Volatility*

The figure presents the percentage drop of the average absolute change in the reciprocal of volatility for each of the 75 futures contracts of the dataset when switching from the standard deviation of past returns (STDEV) volatility estimator to the Yang and Zhang (2000) estimator (YZ). The specific sample period of each contract is reported in Table I.

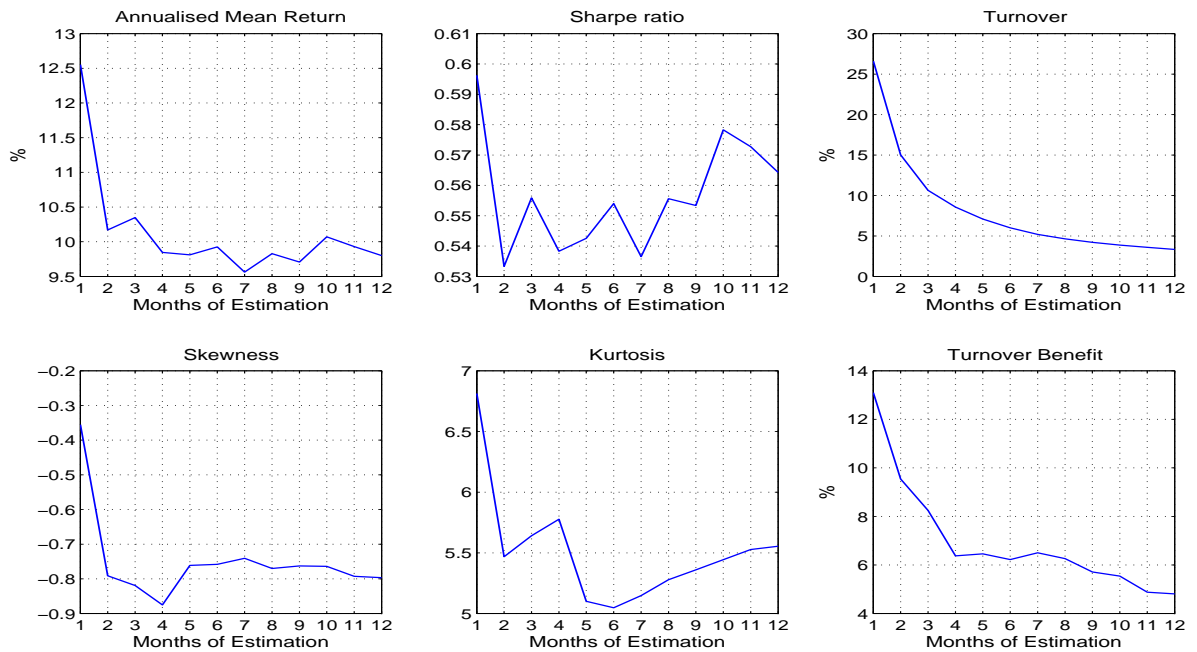


Figure 2: Long-Only Constant Volatility Statistics for Different Estimation Periods

The figure presents the annualised mean return, the Sharpe ratio, the monthly turnover, the skewness and the kurtosis of a long-only constant volatility strategy using Yang and Zhang (2000) volatility estimates across various estimation periods ranging between one to twelve past months. Additionally, the relative turnover benefit for switching from the standard deviation of past returns (STDEV) volatility estimator to the Yang and Zhang (2000) estimator (this turnover benefit denotes a drop in the turnover, but is presented as a positive number) is also presented.




Figure 3: *Time-Series Momentum Statistics for Different Estimation Periods*

The figure presents the annualised mean return, the Sharpe ratio, the monthly turnover, the skewness and the kurtosis of a time-series momentum strategy using Yang and Zhang (2000) volatility estimates across various estimation periods ranging between one to twelve past months. Additionally, the relative turnover benefit for switching from the standard deviation of past returns (STDEV) volatility estimator to the Yang and Zhang (2000) estimator (this turnover benefit denotes a drop in the turnover, but is presented as a positive number) is also presented.

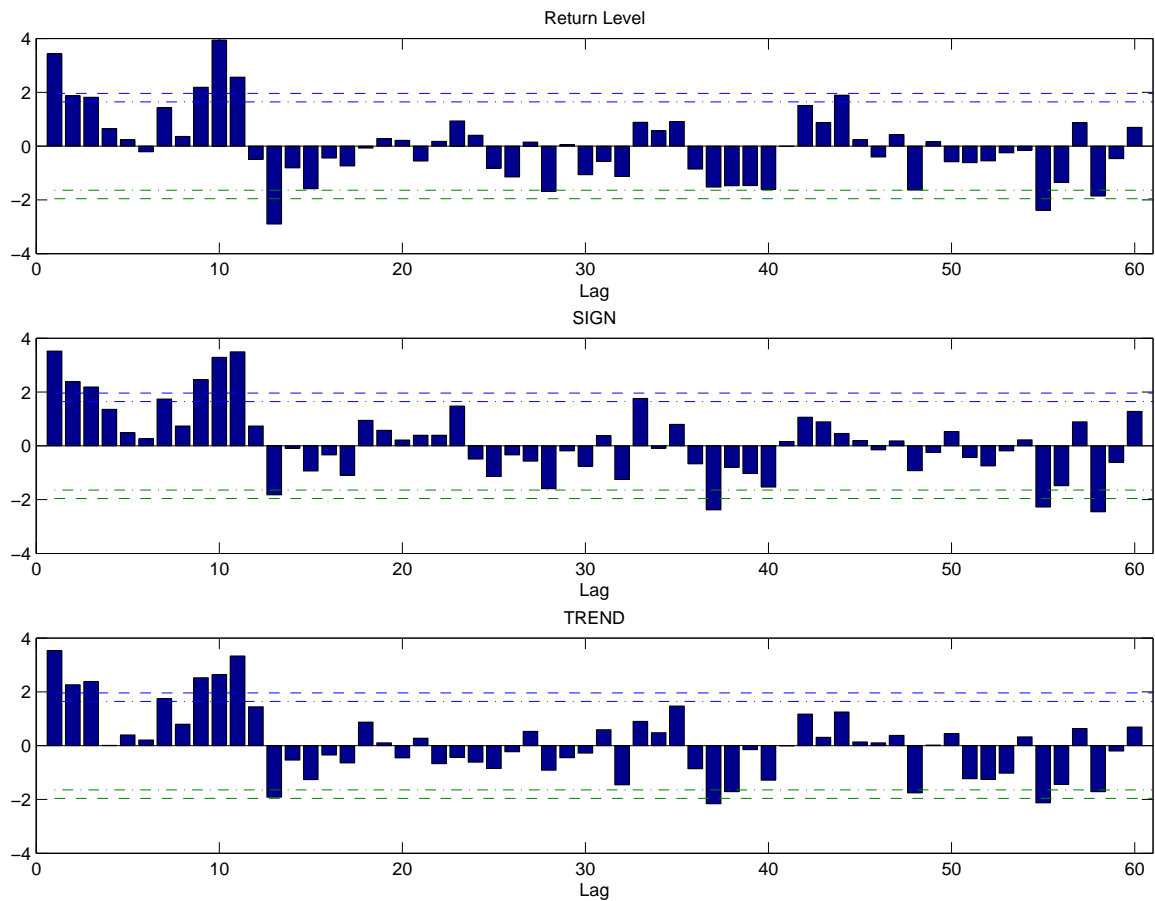


Figure 4: *Time-Series Return Predictability*

The figure presents the t -statistics of the pooled regression coefficient from regressing monthly excess returns of the futures contracts on lagged excess returns or lagged excess momentum trading rules. Panel A presents the results when lagged excess returns are used as the regressor, Panel B when the regressor is the lagged SIGN rule and Panel C when the regressor is the lagged TREND rule. The t -statistics are computed using standard errors clustered by asset and time (Cameron, Gelbach and Miller 2011, Thompson 2011). The volatility estimates are computed using the Yang and Zhang (2000) estimator on a one-month rolling window. The dashed lines represent significance at the 5% and 10% level. The dataset covers the period December 1974 to February 2013.




Figure 5: *The Effect of Sparse Trading Rule*

Panel A presents annualised Sharpe ratios for univariate time-series momentum strategies with 40% target volatility that use the SIGN of past return as trading rule. Additionally, the change in the Sharpe ratio from applying the TREND sparse trading rule is also presented. Panel B presents the percentage drop in the turnover of each univariate strategy when switching between SIGN and TREND momentum trading rules. The volatility estimator that is used across all strategies is the Yang and Zhang (2000) estimator with an estimation period of three months. The specific sample period of each contract is reported in Table I.

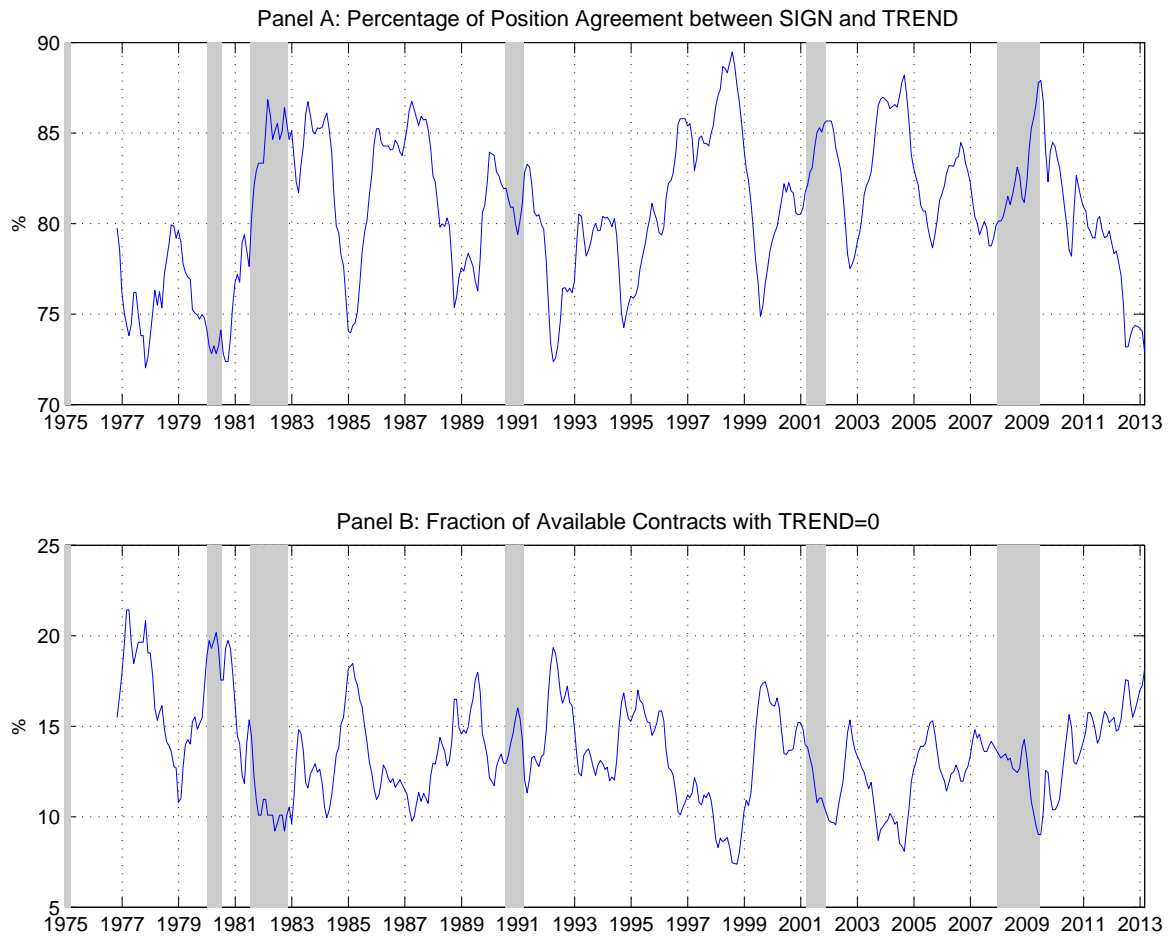


Figure 6: Comparison between *SIGN* and *TREND* Rules

Panel A presents the 12-month moving average of the percentage of contracts at the end of each month for which *SIGN* and *TREND* rules agree (i.e. both long or short for each and every contract). Panel B presents the 12-month moving average of the percentage of available contracts at the end of each month for which the *TREND* rule does not identify a significant upward or downward trend and is therefore equal to zero. The lookback period for which the rules are generated is 12 months and the sample period is December 1975 (first observation in December 1976 due to the 12-month moving average) to February 2013.

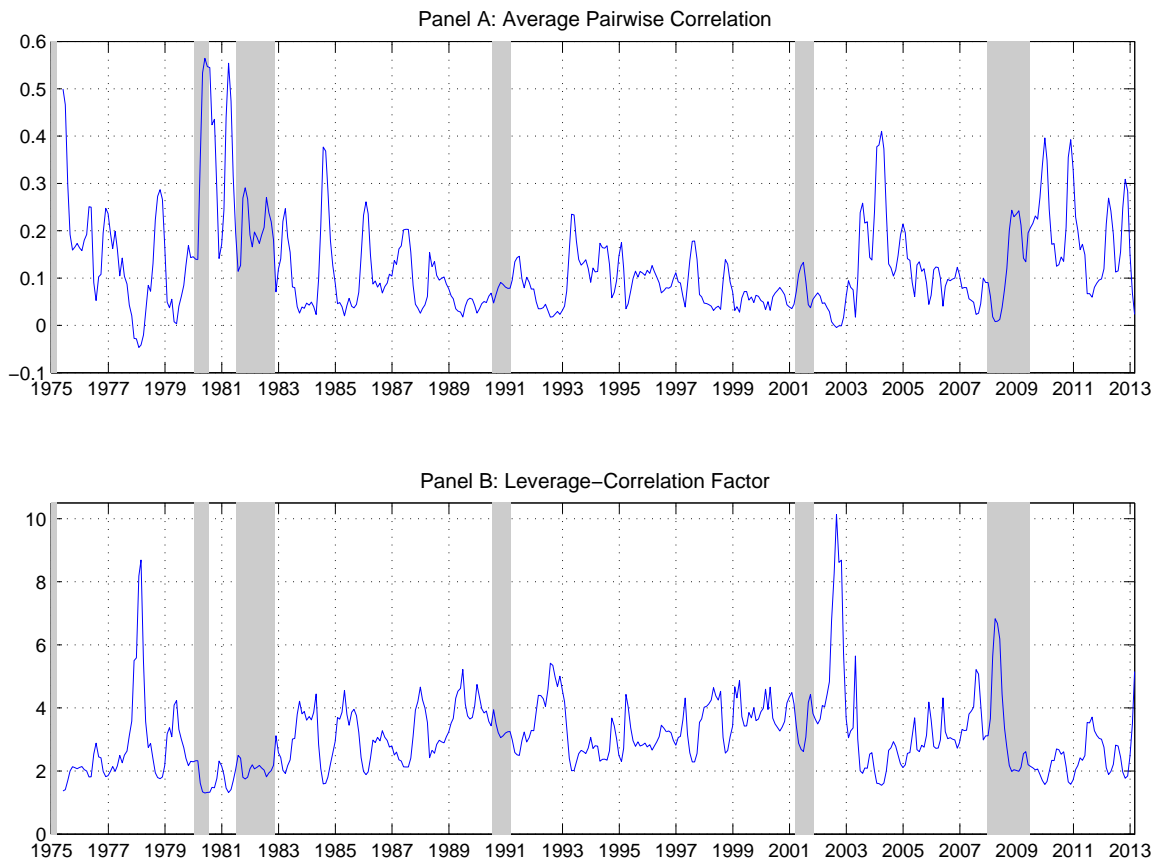


Figure 7: *Average Pairwise Correlation and Correlation Factor*

Panel A presents the 6-month average pairwise correlation of the available contracts at the end of each month. Panel B presents the correlation factor that is deduced from the average pairwise correlation. The sample period is from February 1975 to February 2013.




Figure 8: *The Effect of Correlation in Risk-Adjusted Performance*

The bar chart presents three risk-adjusted performance ratios for the standard time-series momentum strategy and the correlation-adjusted strategy over the entire sample period, December 1975 to February 2013 (Panel A) and over the most recent period following the financial crisis, from January 2009 to February 2013 (Panel B). The volatility estimator that is used is the Yang and Zhang (2000) estimator with an estimation period of three months. The average pairwise correlation is also estimated using a window of three months. The performance ratios are: annualised Sharpe ratio, annualised Sortino ratio (defined as the ratio of the mean return and the downside volatility) and Calmar ratio (defined as the ratio of the mean return and the maximum drawdown).

	Exchange	From	Mean	t-stat	Vol.	Skew	Kurt.	SR
<u>CURRENCIES</u>								
AUD/USD	CME	Feb-1987	5.07	2.11	11.67	-0.40	4.94	0.44
CAD/USD	CME	Feb-1977	0.90	0.79	6.88	-0.31	8.12	0.13
CHF/USD	CME	Dec-1974	0.77	0.36	12.55	0.05	3.81	0.06
EUR/USD	CME	Dec-1974	0.58	0.30	11.37	-0.08	3.55	0.05
GBP/USD	CME	Oct-1977	1.69	0.85	10.70	0.04	4.96	0.16
JPY/USD	CME	Apr-1977	0.69	0.31	11.97	0.49	4.46	0.06
Dollar Index	ICE	Aug-1989	-1.59	-0.84	8.75	0.43	3.83	-0.18
<u>EQUITIES</u>								
DJIA	CBOT	Dec-1974	4.89	1.92	15.22	-0.47	5.37	0.32
NASDAQ 100	CME	Feb-1983	9.29	1.91	25.40	-0.31	4.29	0.37
NYSE Composite	ICE	Dec-1974*	5.19	1.96	15.30	-0.58	5.22	0.34
S&P 500	CME	Dec-1974	5.76	2.20	15.35	-0.48	4.62	0.38
S&P 400 MidCap	CME	Jul-1991	9.59	2.48	17.17	-0.69	5.14	0.56
Russell 2000	ICE	Feb-1988	7.06	1.75	19.17	-0.49	4.01	0.37
DJ Stoxx 50	Eurex	Jan-1987	3.74	1.01	16.49	-0.86	4.98	0.23
Eurostoxx 50	Eurex	Jan-1987	3.88	0.94	18.92	-0.64	4.38	0.21
FTSE 100	NYSE Liffe	Feb-1978	4.41	1.75	16.05	-0.76	5.76	0.28
DAX	Eurex	Dec-1974	4.64	1.33	20.07	-0.48	5.00	0.23
CAC 40	NYSE Liffe	Aug-1987	3.88	0.88	20.64	-0.32	4.12	0.19
IBEX 35	MEFF	Feb-1987	4.92	1.10	22.23	-0.47	4.89	0.22
AEX	NYSE Liffe	Feb-1983	4.81	1.19	20.33	-0.73	5.38	0.24
SMI	Eurex	Aug-1988	6.18	1.66	16.68	-0.56	4.27	0.37
MIB 30	BI	Jan-1998	-0.56	-0.09	22.87	0.00	3.80	-0.02
S&P Canada 60	MX	Feb-1982	4.39	1.38	15.75	-0.68	5.85	0.28
Nikkei 225	CME	Dec-1974	0.60	0.19	19.45	-0.22	4.20	0.03
TOPIX	TSE	Dec-1974	0.37	0.11	17.86	-0.18	4.53	0.02
ASX SPI 200	ASX	Jun-1992*	2.75	0.81	13.58	-0.66	3.70	0.20
Hang Seng	SEHK	Dec-1974	13.01	2.82	28.73	-0.26	5.71	0.45
KOSPI 200	KRX	Feb-1990*	6.51	0.93	31.63	0.86	6.86	0.21
MSCI Taiwan	SGX	Jan-1988	10.28	1.36	34.76	0.43	4.85	0.30
MSCI EAFE	NYSE Liffe	Dec-1974	2.97	1.03	15.90	-0.58	5.32	0.19
<u>INTEREST RATES</u>								
US Treasury Bills 3Mo	CME	Feb-1982*	1.01	2.91	1.44	1.14	8.96	0.70
US Treasury Note 2Yr	CBOT	Feb-1991	1.65	3.73	1.75	0.28	3.44	0.95
US Treasury Note 5Yr	CBOT	Aug-1988	3.23	3.56	4.23	0.05	3.66	0.76
US Treasury Note 10Yr	CBOT	Feb-1983	4.82	3.77	6.90	0.15	3.98	0.70
US Treasury Bond 30Yr	CBOT	Nov-1982	5.93	3.18	10.55	0.25	4.46	0.56
Municipal Bonds	CBOT	Jul-1985*	5.57	3.30	8.04	-0.58	4.62	0.69
Eurodollar 3Mo	CME	Jan-1982	1.19	3.61	1.52	1.09	8.26	0.78
Euribor 3Mo	NYSE Liffe	Feb-1999	0.41	1.56	0.73	1.27	11.72	0.57
Euro/German Schatz 2Yr	Eurex	Apr-1997	1.00	2.43	1.39	0.08	3.59	0.72
Euro/German Bobl 5Yr	Eurex	Feb-1997	2.71	2.94	3.29	-0.02	2.70	0.83
Euro/German Bund 10Yr	Eurex	Feb-1997	4.17	2.98	5.35	0.08	2.88	0.78
Euro/German Buxl 30Yr	Eurex	Oct-2005	5.53	1.28	12.64	1.02	4.83	0.44
Australian 3Yr	ASX	Aug-2001	0.47	1.25	1.08	0.45	2.88	0.43
Australian 10Yr	ASX	Aug-2001	0.36	1.26	0.93	0.27	2.96	0.38
UK Sterling 3Mo	NYSE Liffe	Aug-1998	0.75	2.33	0.92	1.72	12.55	0.82
UK Long Gilt	NYSE Liffe	Aug-1998	2.96	1.82	5.97	0.28	3.59	0.50
Canadian 10Yr	MX	May-1990	4.76	3.88	5.87	-0.04	3.23	0.81
Japanese 10Yr	TSE	Aug-2003	1.75	2.02	2.99	-0.73	4.99	0.59
Korean 3Yr	KRX	Sep-2003*	1.69	1.63	3.08	0.88	6.64	0.55

(Continued on next page)

(Continued from previous page)

	Exchange	From	Mean	t-stat	Vol.	Skew	Kurt.	SR
COMMODITIES								
<u>ENERGY</u>								
Light Crude Oil	NYMEX	Feb-1987	13.22	1.71	34.13	0.40	5.46	0.39
Brent Crude Oil	NYMEX	Sep-2003	15.45	1.22	31.37	-0.63	4.81	0.49
Heating Oil	NYMEX	Feb-1984	13.53	2.09	33.75	0.49	4.84	0.40
Natural Gas	NYMEX	Feb-1993	0.01	0.00	61.39	1.03	5.77	0.00
RBOB Gasoline	NYMEX	Oct-1987	22.10	2.99	36.24	0.36	5.45	0.61
<u>METALS</u>								
Copper	COMEX	Jan-1990	9.47	1.48	26.45	-0.05	5.46	0.36
Gold	COMEX	Feb-1984	1.88	0.72	15.42	0.31	4.12	0.12
Palladium	NYMEX	Feb-1994	14.37	1.60	35.29	0.34	5.68	0.41
Platinum	NYMEX	Aug-2003	11.83	1.19	27.03	-0.81	7.33	0.44
Silver	COMEX	Jan-1984	2.71	0.56	27.80	0.26	4.17	0.10
<u>MEAT</u>								
Feeder Cattle	CME	Feb-1978	2.46	1.04	14.52	-0.38	5.20	0.17
Live Cattle	CME	Dec-1974	4.66	1.70	16.51	-0.13	4.39	0.28
Live Hogs	CME	Dec-1974	3.32	0.83	25.57	-0.04	3.31	0.13
Pork Bellies	CME	Dec-1974*	0.80	0.15	36.87	0.44	4.23	0.02
<u>GRAINS</u>								
Corn	CBOT	Aug-1982	-0.92	-0.19	25.95	0.64	6.02	-0.04
Oats	CBOT	Aug-1982	-1.02	-0.16	34.68	2.76	25.57	-0.03
Soybean Oil	CBOT	Aug-1982	2.26	0.48	26.29	0.58	6.13	0.09
Soybean Meal	CBOT	Aug-1982	8.70	1.90	25.19	0.24	3.88	0.35
Soybeans	CBOT	Aug-1982	4.43	1.08	23.53	0.13	4.13	0.19
Wheat	CBOT	Aug-1982	-2.98	-0.68	25.47	0.36	4.92	-0.12
<u>SOFTS</u>								
Cocoa	ICE	Aug-1986	-3.81	-0.76	29.17	0.58	4.14	-0.13
Coffee	ICE	Feb-1987	-2.04	-0.26	37.86	1.00	5.69	-0.05
Cotton	ICE	Feb-1987	1.42	0.24	26.01	0.28	3.78	0.05
Lumber	CME	Dec-1974	-3.43	-0.68	29.17	0.29	3.67	-0.12
Orange Juice	ICE	Aug-1987	2.87	0.47	32.25	0.68	4.57	0.09
Sugar	ICE	Aug-1986	8.77	1.34	33.10	0.33	3.81	0.27

Table I: Summary Statistics for Futures Contracts

The table presents summary statistics for the 75 futures contracts of the dataset, which are estimated using monthly fully collateralised excess return series. The statistics are: annualised mean return in %, Newey and West (1987) t-statistic, annualised volatility in %, skewness, kurtosis and annualised Sharpe ratio (SR). The table also indicates the exchange that each contract is traded at the end of the sample period as well as the starting month and year for each contract. All but 7 contracts have data up until February 2013. The remaining 7 contracts are indicated by an asterisk (*) next to the starting date and their sample ends prior to February 2013: NYSE Composite up to January 2012, ASX SPI 200 up to January 2012, KOSPI 200 up to January 2012, US Treasury Bills 3Mo up to August 2003, Municipal Bonds up to March 2006, Korean 3Yr up to June 2011 and Pork Bellies up to April 2011. The EUR/USD contract is spliced with the DEM/USD (Deutsche Mark) contract for dates prior to January 1999 and the RBOB Gasoline contract is spliced with the Unleaded Gasoline contract for dates prior to January 2007, following Moskowitz, Ooi and Pedersen (2012).

Panel A: Long-Only Constant Volatility Strategies			
	STDEV	YZ	PF
Mean (%)	12.69	12.55	14.69
Volatility (%)	20.98	21.08	16.50
Skewness	-0.30	-0.36	-0.41
Kurtosis	6.49	6.81	3.55
CAPM Beta	0.97 (12.36)	0.99 (12.55)	0.80 (11.93)
Sharpe Ratio	0.61	0.60	0.89
LW p-value(%)	0.10	0.08	H_0
Monthly Turnover (%)	30.63	26.61	29.44
Turnover benefit vs. STDEV (%)	0.00	-13.10	-3.87
Panel B: Time-Series Momentum Strategies			
	STDEV	YZ	PF
Mean (%)	14.95	14.72	17.33
Volatility (%)	17.96	17.91	13.34
Skewness	-2.38	-2.31	-0.03
Kurtosis	27.88	26.79	3.02
CAPM Beta	0.07 (0.66)	0.08 (0.68)	0.02 (0.21)
Sharpe Ratio	0.83	0.82	1.30
LW p-value(%)	0.08	0.09	H_0
Monthly Turnover (%)	46.51	42.79	45.52
Turnover benefit vs. STDEV (%)	0.00	-8.00	-2.13

Table II: *The Effect of Volatility Estimator*

The table presents performance statistics for long-only constant volatility strategies (Panel A) and time-series momentum strategies (Panel B) that differ between each other in the volatility estimator used: standard deviation of past returns (STDEV) or the Yang and Zhang (2000) estimator (YZ). The ex-ante volatility estimation period is one month. For comparison purposes, the last column reports statistics for a strategy that uses the ex-post realised volatility over the holding period, i.e. the Perfect Foresight estimator (PF). The reported statistics are: annualised mean return in %, annualised volatility in %, skewness, kurtosis, CAPM beta with the respective Newey and West (1987) t-statistic, annualised Sharpe ratio, Ledoit and Wolf (2008) p-value for the null hypothesis of equality of Sharpe ratios between all different strategies with the PF strategy, monthly turnover in % and relative turnover benefit from switching between STDEV estimator and any other volatility estimator. The dataset covers the period December 1975 to February 2013.

	SIGN	TREND
Mean (%)	15.28	14.83
Volatility (%)	14.74	14.96
Skewness	-0.20	-0.28
Kurtosis	3.99	3.86
CAPM Beta	0.05 (0.45)	0.08 (0.79)
Sharpe Ratio	1.04	0.99
LW p-value(%)	53.31	
Monthly Turnover (%)	28.89	9.68
Turnover benefit (%)	0.00	-66.49
Correlation	0.92	

Table III: *Time-series Momentum Strategies and the Effect of Trading Rule*

The table presents performance statistics for the time-series momentum strategies that differ between each other in the momentum trading rule used: sign of past return (SIGN) versus the t-statistic of a linear trend fit on the price path (TREND). The volatility estimator that is used is the Yang and Zhang (2000) estimator with an estimation period of three months. The reported statistics are: annualised mean return in %, annualised volatility in %, skewness, kurtosis, CAPM beta with the respective Newey and West (1987) t-statistic, annualised Sharpe ratio, Ledoit and Wolf (2008) p-value for the null hypothesis of equality of Sharpe ratios, monthly turnover in %, relative turnover from switching between SIGN to TREND rule and finally correlation between the two strategies. The dataset covers the period December 1975 to February 2013.

Panel A: Full-Sample: December 1975 - February 2013		
	Unadjusted	Correlation-Adjusted
Mean (%)	15.28	15.98
Volatility (%)	14.74	15.17
Skewness	-0.20	0.32
Kurtosis	3.99	7.29
Sharpe Ratio (Mean/Volatility)	1.04	1.05
Sortino Ratio (Mean/Downside Vol.)	1.81	1.98
Calmar Ratio (Mean/Max Drawdown)	0.57	0.70
Monthly Turnover (%)	28.89	41.23
Correlation		0.94
Panel B: After Financial Crisis: January 2009 - February 2013		
	Unadjusted	Correlation-Adjusted
Mean (%)	2.00	3.52
Volatility (%)	14.27	12.21
Skewness	-0.07	0.64
Kurtosis	3.69	6.86
Sharpe Ratio (Mean/Volatility)	0.14	0.29
Sortino Ratio (Mean/Downside Volatility)	0.20	0.46
Calmar Ratio (Mean/Max Drawdown)	0.12	0.24
Monthly Turnover (%)	20.33	29.49
Correlation		0.96

Table IV: *Time-Series Momentum Strategies and the Effect of Correlation*

The table presents performance statistics for the standard time-series momentum strategy and the correlation-adjusted strategy. The volatility estimator that is used is the Yang and Zhang (2000) estimator with an estimation period of three months. The average pairwise correlation is also estimated using a window of three months. The reported statistics are: annualised mean return in %, annualised volatility in %, skewness, kurtosis, annualised Sharpe ratio, annualised Sortino ratio (defined as the ratio of the mean return and the downside volatility), Calmar ratio (defined as the ratio of the mean return and the maximum drawdown), monthly turnover in % and finally correlation between the two strategies. Panel A covers the entire sample period December 1975 to February 2013, whereas Panel B covers the most recent period following the financial crisis, from January 2009 to February 2013.

Panel A: The Effect of the Volatility Estimator				
	Long-Only		Time-Series Momentum	
	STDEV	YZ	STDEV	YZ
Mean (%)	12.69	12.55	14.95	14.72
Sharpe Ratio	0.61	0.60	0.83	0.82
Monthly Turnover (%)	30.63	26.61	46.51	42.79
<i>After Costs of 10bps:</i>				
Realised Costs (%)	0.37	0.32	0.56	0.51
Sharpe Ratio	0.59	0.58	0.80	0.79
<i>After Costs of 50bps:</i>				
Realised Costs (%)	1.84	1.60	2.79	2.57
Sharpe Ratio	0.52	0.52	0.68	0.68
Panel B: The Effect of the Trading Rule				
	SIGN	TREND		
Mean (%)	15.28	14.83		
Sharpe Ratio	1.04	0.99		
Monthly Turnover (%)	28.89	9.68		
<i>After Costs of 10bps:</i>				
Realised Costs (%)	0.35	0.12		
Sharpe Ratio	1.01	0.98		
<i>After Costs of 50bps:</i>				
Realised Costs (%)	1.73	0.58		
Sharpe Ratio	0.92	0.95		
Panel C: The Effect of the Correlations				
	<i>Full Sample</i>		<i>2009-2013</i>	
	Unadjusted	Corr.-Adj.	Unadjusted	Corr.-Adj.
Mean (%)	15.28	15.98	2.00	3.52
Sharpe Ratio	1.04	1.05	0.14	0.29
Monthly Turnover (%)	28.89	41.23	20.33	29.49
<i>After Costs of 10bps:</i>				
Realised Costs (%)	0.35	0.49	0.24	0.35
Sharpe Ratio	1.01	1.02	0.12	0.26
<i>After Costs of 50bps:</i>				
Realised Costs (%)	1.73	2.47	1.22	1.77
Sharpe Ratio	0.92	0.89	0.05	0.14

Table V: The Effect of Trading Costs

The table presents the annualised realised costs and the respective after-costs Sharpe ratios for the long-only strategies and time-series momentum strategies that employ different volatility estimators (Panel A), for time-series momentum strategies that employ different trading rules (Panel B) and for correlation-adjusted time-series momentum strategies that employ dynamic leverage as a function of the average pairwise correlation (Panel C). The estimated costs assume two different levels of the average market impact; 10 basis points and 50 basis points. The sample period is December 1975 to February 2013, except for the last two columns of Panel C, which focus on the most recent period following the financial crisis, from January 2009 to February 2013.