Taylor, N. (2017). Timing strategy performance in the crude oil futures market. *Energy Economics*, 66, 480-492.
<https://doi.org/10.1016/j.eneco.2017.07.019>

Peer reviewed version

License (if available):
CC BY-NC-ND

Link to published version (if available):
[10.1016/j.eneco.2017.07.019](https://doi.org/10.1016/j.eneco.2017.07.019)

[Link to publication record in Explore Bristol Research](#)
PDF-document

This is the accepted author manuscript (AAM). The final published version (version of record) is available online via Elsevier at <https://doi.org/10.1016/j.eneco.2017.07.019> . Please refer to any applicable terms of use of the publisher.

University of Bristol - Explore Bristol Research

General rights

This document is made available in accordance with publisher policies. Please cite only the published version using the reference above. Full terms of use are available:
<http://www.bristol.ac.uk/red/research-policy/pure/user-guides/ebr-terms/>

Timing strategy performance in the crude oil futures market

Nick Taylor^a

^a*School of Economics, Finance and Management, University of Bristol, Bristol, BS8 1TN, UK.*

Abstract

The rewards to speculative trading in the crude oil futures market are assessed. For investors who adopt timing strategies that maximise their (iso-elastic) utility during each trading session, the rewards can be economically significant providing that transaction costs are small. Moreover, we are able to show via a decomposition of performance that the bulk of this benefit is due to their ability to predict realised volatility (that is, the second realised moment). The benefits derived from predicting other realised moments either require unrealistic levels of skill (all odd moments) or an infeasible degree of risk aversion (the fourth moment and higher even moments).

Keywords: Crude oil futures, timing strategies, realised moments, volatility.

JEL: C10, C22, C53, C58, G11, G17.

Email address: `nick.taylor@bristol.ac.uk` (Nick Taylor)

1. Introduction

There is sustained interest in speculative trading strategies that take temporal positions in assets (henceforth *timing strategies*). While practitioners tend to focus on promoting (only) strategies that beat the market, the academic community takes a more skeptical stance. The traditional approach of the latter has been to critically evaluate the performance of *market timing* strategies based on forecasts of future returns (that is, first realised moment forecasts); see Kandel and Stambaugh (1996) and Welch and Goyal (2008) for seminal examinations of equity return predictability amongst a huge literature, and Wang and Yang (2010), Kristoufek and Vosvrda (2014), Liu et al. (2015), Lubnau and Todorova (2015), Wang et al. (2016a) and Liu et al. (2017a, 2017b) for recent applications in the context of energy futures trading.¹ We build on this literature by considering the performance of investors who seek to maximise their (iso-elastic) utility during each trading session by taking temporal positions in crude oil futures contracts traded on the Chicago Mercantile Exchange (CME).

The previously documented mixed performance of market timing strategies has led to growing interest in the predictability (or otherwise) of higher order return moments. For instance, a number of studies reveal that the performance of *volatility timing* strategies is generally economically significant; see, e.g., West et al. (1993), Fleming et al. (2001, 2003), Marquering and Verbeek (2004), Chiriac and Voev (2011) and Taylor (2014a) for applications to equity data, and Wang et al. (2016b) and Kang et al. (2017) for recent applications in the context of energy futures trading.² We unify the market and volatility timing literatures by considering a framework in which the performance of more general timing strategies is decomposed into the forecasting ability of all realised moments; Jondeau and Rockinger (2012) refer to similar strategies based on non-realised moments as *distribution timing* strategies. However, given our use of realised moments we refer to ours as *realised distribution timing* (henceforth ReDiT) strategies. The proposed framework is able to identify the drivers of ReDiT strategy performance within the context of crude oil futures trading. This is the primary contribution of the paper.

---

¹Most market timing strategies are binary in nature in the sense that they involve constructing forecasts of broad asset classes in order to take a position in the market or not. Merton (1981) refers to investors who employ these strategies as *macro-forecasters*. This is in contrast to *micro-forecasters* who seek mispriced individual stocks.

²A volatility timing strategy typically involves taking a position in a security based on its predicted volatility.

There are good reasons why higher moments are important within the context of trading strategies. Consider an investor who seeks to maximise her utility during a trading session. Providing that the utility function depends on the strategy returns and is infinitely differentiable (at a real or complex number), then it follows that a Taylor series expansion can be applied to give a function that is linear in terms of all return moments. A seminal example of this approach is Levy and Markowitz (1979) who consider the performance of an approximation to expected utility via the first two return moments; see Garlappi and Skoulakis (2011) for details of the properties of this approximation and those based on inclusion of higher moments. Consequently, the performance of a timing strategy that seeks to maximise such a function is determined by the ability to forecast each of the return moments. This is the underlying approach in the current paper.

Under a set of realistic assumptions (including transaction costs), we are able to provide an explicit expression for the performance of optimal ReDiT strategies as a function of a simple measure of ability to forecast individual realised (return) moments. Using this expression we consider the benefits of trading futures contracts in a leading energy futures market, viz., crude oil. For utility maximising investors, the specific conditions under which there are benefits to ReDiT strategies are identified. The results indicate that the benefits to employing these strategies rest solely on investors' ability to forecast the second realised moment (that is, realised volatility). This is because either investors have little ability to forecast the other realised moments (the odd realised moments) or require a huge degree of risk aversion (the fourth realised moment). Transaction costs are also an important consideration. Only when transactions are small are benefits available. The results reinforce the efficient nature of energy futures markets.

The rest of the paper is organised as follows. The next section provides the investment framework that includes an explicit expression for the expected performance of ReDiT strategies. This section also includes a description of the models used to generate forecasts of the realised moments. Section 3 contains the application to crude oil futures data, and the final section concludes.

2. Methodologies

This section contains the framework within which investors are assumed to operate, and a description of the models used by these investors to generate forecasts of the realised (return) moments.

2.1. The investment framework

Investors accord to the following set of assumptions.

Assumption 1. Trading takes place during the t th trading session. There are S intra-session periods in each session such that $s = 1, \dots, S$. Intra-session returns to the ReDiT strategy are denoted $R_{k,s,t}$, where the subscripts represent the k th ReDiT strategy associated with the s th intra-session period of the t th session.

Assumption 2. Investor utility within each session (henceforth *intra-session utility*) belongs to the iso-elastic (power) utility function given by

$$f[R_{k,s,t}] = \frac{(1 + R_{k,s,t})^{1-\gamma} - 1}{1 - \gamma}, \quad (1)$$

where $\gamma \geq 0$ and $R_{k,s,t} > -1$. Note that minus unity is included in the numerator so that all investors (irrespective of their risk preferences, that is, their γ value) have zero utility when the return to the strategy is zero.

Remark. This utility function can be approximated as an N th order polynomial series. To minimise the approximation error associated with this series we take a log transformation of returns such that $r_{k,s,t} = \ln[1 + R_{k,s,t}]$. Taking this expansion about κ we obtain

$$\tilde{f}[R_{k,s,t}] = \sum_{n=0}^N \beta_n \delta_{k,s,t}^n, \quad (2)$$

where

$$\beta_n = \begin{cases} (1 - \gamma)^{-1}(e^{(1-\gamma)\kappa} - 1), & \text{if } n = 0, \\ (1 - \gamma)^{n-1}e^{(1-\gamma)\kappa}/n!, & \text{otherwise.} \end{cases} \quad (3)$$

Here $\delta_{k,s,t} = r_{k,s,t} - \kappa$ is the stochastic deviation.

Assumption 3. Over each session, ReDiT strategies involve either investing all wealth in a single risky asset or all wealth in a safe asset earning zero interest. This decision is determined in the previous session and is based on an information set denoted $\mathcal{F}_{k,t-1}$.

Remark. The binary nature of the ReDiT strategies is consistent with the market timing strategies proposed and analysed previously; see Merton (1981).

Assumption 4. No inventory is held in the period between trading sessions. That is, all (open) positions are closed out at the end of each session.

Remark. This assumption is the hallmark of modern investment strategies such the low-latency trading strategies employed by high-frequency traders; see Jones (2013) for an overview of this literature.

Assumption 5. Each ReDiT strategy is subject to a transaction cost, incurred whenever a trade in the risky asset occurs.

Remark. Assumptions 4 and 5 together imply that returns to each ReDiT strategy are given by

$$1 + R_{k,s,t} = (1 + x_{k,t-1}R_{s,t})(1 - c_{k,t-1}), \quad (4)$$

where $R_{s,t}$ is the return to the risky asset, $x_{k,t-1}$ is the trade indicator such that it equals unity (trade) or zero (no trade), $c_{k,t-1} = \tau x_{k,t-1}$ is the total cost of trading, and τ is the transaction cost associated with trading the asset.

Remark. Taking logs of (4) and subtracting κ we obtain an expression in terms of the stochastic deviation, that is,

$$\delta_{k,s,t} = \ln[1 + x_{k,t-1}R_{s,t}] + \ln[1 - c_{k,t-1}] - \kappa, \quad (5)$$

where previous notation is maintained.

Assumption 6. Each ReDiT strategy seeks maximisation of the summation of utility during the t th session. This summation is given by

$$U_{k,t} = \sum_{s=1}^S \tilde{f}[R_{k,s,t}] = \sum_{n=0}^N \beta_n \sum_{s=1}^S \delta_{k,s,t}^n. \quad (6)$$

This is henceforth referred to as *realised session utility*.

Remark. Using the expression in (5), substituting into (6), and applying the binomial theorem we

obtain

$$\begin{aligned}
U_{k,t} &= \sum_{n=0}^N \beta_n \sum_{s=1}^S (\ln[1 + x_{k,t-1} R_{s,t}] + \ln[1 - c_{k,t-1}] - \kappa)^n, \\
&= \sum_{n=0}^N \beta_n \sum_{s=1}^S \sum_{m=0}^n \binom{n}{m} (\ln[1 - c_{k,t-1}])^{n-m} (\ln[1 + x_{k,t-1} R_{s,t}] - \kappa)^m.
\end{aligned} \tag{7}$$

The advantage of this decomposition will become apparent in the subsequent analysis when optimal ReDiT strategy behavior is considered.

Assumption 7. To ensure that each ReDiT strategy is time consistent, investors maximise conditional expectations of (7). Hence their objective function is given by

$$\max_{x_{k,t-1}} E[U_{k,t} | \mathcal{F}_{k,t-1}] = \max_{x_{k,t-1}} \sum_{n=0}^N \beta_n \sum_{s=1}^S \sum_{m=0}^n \binom{n}{m} (\ln[1 - c_{k,t-1}])^{n-m} E[(\ln[1 + x_{k,t-1} R_{s,t}] - \kappa)^m | \mathcal{F}_{k,t-1}], \tag{8}$$

where previous notation is maintained.

Assumption 8. Individual realised moments and their respective (unbiased) forecasts are related as follows:

$$y_{n,t} = E[y_{n,t} | \mathcal{F}_{k,t-1}] + \epsilon_{k,n,t}, \tag{9}$$

where $\epsilon_{k,n,t}$ is the forecast error. Furthermore, $E[y_{n,t} | \mathcal{F}_{k,t-1}]$ and $\epsilon_{k,n,t}$ are independent such that $cov[E[y_{n,t} | \mathcal{F}_{k,t-1}], \epsilon_{k,n,t}] = 0$.

Remark. The quality of the individual realised moment forecasts is measured by the coefficient of determination statistic (denoted $\mathcal{R}_{k,n}^2$) associated with (9). This statistic represents our measure of *investor skill* with respect to the k th strategy and the n th realised moment. It follows that the variance of $E[y_{n,t} | \mathcal{F}_{k,t-1}]$ is given by $\sigma_n^2 \mathcal{R}_{k,n}^2$, where σ_n^2 denotes the variance of $y_{n,t}$.

Proposition 1. *Under Assumptions 1 to 8, the unconditional expectation of the realised session utility to the k th ReDiT strategy is given by*

$$E[U_{k,t}] = \int_{w=0}^{\infty} w g_{E[z_{1,t} | \mathcal{F}_{k,t-1}]}[w] dw, \tag{10}$$

where $g_{E[z_{1,t} | \mathcal{F}_{k,t-1}]}[\cdot]$ is the probability density function (PDF) associated with $E[z_{1,t} | \mathcal{F}_{k,t-1}]$, which

in turn is given by

$$\mathbb{E}[z_{1,t}|\mathcal{F}_{k,t-1}] = \sum_{n=0}^N \beta_n \sum_{m=0}^n \binom{n}{m} (\ln[1 - \tau])^{n-m} \mathbb{E}[y_{m,t}|\mathcal{F}_{k,t-1}]. \quad (11)$$

Here $z_{1,t}$ is the realised session utility obtained when a trade occurs, and

$$y_{m,t} = \sum_{s=1}^S (\ln[1 + R_{s,t}] - \kappa)^m \quad (12)$$

is the m th realised moment associated with returns to the risky asset being traded.

Proof. See appendix.

Remark. This proposition provides an expression for a measure of strategy performance given by the unconditional expectation of realised session utility. This is shown to be a function of the distribution of the conditional expectation of realised session utility associated with trading. This conditional expectation is, in turn, a function of the conditional expectations of the realised moments. It is noticeable that the truncated nature of the integral in (10) means that the higher the variance of these conditional expectations, given by $\sigma_n^2 \mathcal{R}_{k,n}^2$ the greater the performance of the strategy (ceteris paribus). It follows that investor skill (given by $\mathcal{R}_{k,n}^2$) and performance are positively related.

To obtain the conditional expectations of the realised moments we require models of these moments. These are described in the next section.

2.2. Expectation formation

Conditional expectations of realised moments are based on the assumption that the n th realised moment evolves according to the following fractionally integrated moving average (FIMA) process:

$$(1 - L)^d y_{n,t+1} = (1 - \theta L) \xi_{t+1}, \quad (13)$$

where L is the lag operator, d is the fractional order of differencing, and ξ_{t+1} is a suitably defined error term; see Proietti (2016) for a recent application of the FIMA model in the context of realised

moment modeling.³ It is useful to write (13) as

$$y_{n,t+1} - g[d, L]y_{n,t} = (1 - \theta L)\xi_{t+1}, \quad (14)$$

where

$$g[d, L] = \sum_{k=1}^{\infty} \binom{d}{k} (-L)^{k-1} = d + \frac{1}{2}d(1-d)L + \frac{1}{6}d(1-d)(2-d)L^2 + \dots, \quad (15)$$

with previous notation maintained.

Taking conditional expectations of (13) we obtain

$$\mathbb{E}[y_{n,t+1}|\mathcal{F}_t] = g[d, L]y_{n,t} - \theta\xi_t. \quad (16)$$

Noting that $y_{n,t} = \mathbb{E}[y_{n,t}|\mathcal{F}_{t-1}] + \xi_t$ and rearranging gives

$$\mathbb{E}[y_{n,t+1}|\mathcal{F}_t] = (g[d, L] - \theta)y_{n,t} + \theta\mathbb{E}[y_{n,t}|\mathcal{F}_{t-1}]. \quad (17)$$

Thus the next period expectation is a function of current (and past) realisations, and the current expectation. These expectations are henceforth referred to as *long-memory adaptive expectations*.

The expression in (17) encompasses other expectation formation schemes. In particular, if we set d and θ equal to zero, then we obtain

$$\mathbb{E}[y_{n,t+1}|\mathcal{F}_t] = 0. \quad (18)$$

These are henceforth referred to as *naive expectations*. Alternatively, if d is set to unity then we obtain the following:

$$\mathbb{E}[y_{n,t+1}|\mathcal{F}_t] = (1 - \theta)y_{n,t} + \theta\mathbb{E}[y_{n,t}|\mathcal{F}_{t-1}]. \quad (19)$$

These coincide with the commonly used *adaptive expectations*.

---

³Intercepts have been suppressed for presentational convenience only.

3. Results

This section contains a description (including summary statistics) of the data used, and examines the performance of the ReDiT strategies within the context of trading crude oil futures contracts during trading sessions of various length.

3.1. Data

Realised moments are constructed using intraday data observed during the daytime trading session. In particular, prices associated with all trades in all futures contracts on crude oil (CME traded with CL ticker) collected over the period January 1, 1988 to December 31, 2015. These data were obtained from Tickdata, Inc. These transaction data are firstly converted to five minute frequency data, then log price differences are calculated with all overnight price differences eliminated.⁴ Finally, a single continuous price difference series is constructed by only using data pertaining to the nearest maturity contract. The first six realised moments are constructed by summing powers of this series over an intra-daily (defined as summing over the morning and afternoon trading sessions separately), daily, weekly and bi-weekly frequency.⁵ Use of data over these frequencies corresponds to investors assuming intra-daily, daily, weekly and bi-weekly trading sessions.

3.2. Estimation details

The following sections describe how we estimate the expectations of investors, and how the result in Proposition 1 is implemented.

3.2.1. Estimating expectations

Naive, adaptive and long-memory adaptive expectations are formed using the models described in section 2.2. The underlying FIMA model (and restrictions thereof) is estimated using each realised moment observed during each trading session. Estimation is achieved by minimising the sum of squared 1-step ahead prediction errors using the Newton-Raphson and cubic/quadratic step length methods, implemented via the CO package in GAUSS v.17. To ensure time consistency, the

---

⁴Construction of realised moments using five minute frequency data is standard practice. See Liu et al. (2015) for the virtues of using this frequency in the context of constructing realised volatility.

⁵The subsequent analysis only considers the first six realised moments. As will be revealed there is little purpose in considering higher order realised moments as predictability of these will not add any value to ReDiT strategy users. Indeed, analysis of the first four moments is sufficient to reach this conclusion.

models are estimated using a rolling window of observations. This scheme adds one observation to the end of the estimation sample and removes one observation from the start of the sample until the December 31, 2015 observation is reached.⁶ At each point in time out-of-sample conditional expectations of each realised moment are constructed. These expectations cover the period January 1, 1992 to December 31, 2015.

3.2.2. Estimating expected performance

To operationalise the result in Proposition 1 we adopt the following procedure. We have broken this down into the following steps.

1. The estimated forecast errors associated with each model, denoted $\hat{\epsilon}_{k,n,t}$ in (9), are calculated for each the above expectation schemes (that is, naive, adaptive, and long-memory adaptive expectations).
2. The estimated out-of-sample $\mathcal{R}_{k,n}^2$ values (denoted $\hat{\mathcal{R}}_{k,n}^2$) associated with each moment are calculated using the estimated forecast errors calculated in the previous step.
3. The realised moment expectations are calibrated such that they have a variance equal to $\hat{\sigma}_n^2 \hat{\mathcal{R}}_{k,n}^2$, where $\hat{\sigma}_n^2$ is the estimated variance of the n th realised moment. These conditional expectations are combined using (11) to give the conditional expectations of $z_{1,t}$, that is, $\hat{\mathbb{E}}[z_{1,t}|\mathcal{F}_{k,t-1}]$. Note that should one wish to consider strategies based on alternative skill levels then this can be accommodated by selected the desired $\mathcal{R}_{k,n}^2$ value in this step.
4. The PDF associated with $\hat{\mathbb{E}}[z_{1,t}|\mathcal{F}_{k,t-1}]$ (that is, $g_{\hat{\mathbb{E}}[z_{1,t}|\mathcal{F}_{k,t-1}]}[\cdot]$) is calculated using a non-parametric kernel estimate based on the Epanechnikov kernel window. An automatic smoothing parameter value, denoted h_{opt} , is used and given by $h_{opt} = 0.9\hat{A}n^{-1/5}$, where $\hat{A}$ is the minimum of the standard deviation of $\hat{\mathbb{E}}[z_{1,t}|\mathcal{F}_{k,t-1}]$ and the interquartile range of these expectations divided by 1.34, and n is the number of observations.⁷
5. The integral in Proposition 1 involving $g_{\hat{\mathbb{E}}[z_{1,t}|\mathcal{F}_{k,t-1}]}[\cdot]$ is calculated numerically by summing the probability-weighted values of $\hat{\mathbb{E}}[z_{1,t}|\mathcal{F}_{k,t-1}]$ in the positive region.

---

⁶We adopt a 1000-day rolling window.

⁷This choice of smoothing parameter “will do very well for a wide range of densities” (Silverman, 1986).

3.3. The trading parameters

Aside from transaction costs, there are two sets of parameters that will determine the performance of the ReDiT strategies: those related to market conditions and those related to investor skill. These are described in the following sections.

3.3.1. Realised moments

To assess the nature of the former we present the mean, standard deviation, skewness and kurtosis (excess) of the first six realised moments. These are calculated over the intra-daily, daily, weekly and bi-weekly trading sessions and presented in Table 1.

Insert Table 1 here

The mean of the first moment over the daily trading session implies an annualised mean (log) return of 8.3%. This is associated with a mean second moment which corresponds to an annualised return volatility of 27.4%. The results also show that the distribution of realised moments is (unsurprisingly) highly non-normal over all trading session lengths. For instance, the excess kurtosis values associated with the first realised moment over the daily trading session is 3.6. This is an important factor regarding the use of the formulae derived in the previous section.

The non-normality of the realised moments most likely means that the formula in Proposition 2 in the appendix to this paper cannot be used. This is because this formula relies of the conditional expectations of $z_{1,t}$ being normally distributed, which in turn are a linear combination of the expectations of the realised moments; see equation (13). To examine this issue more closely we provide plots of the densities of the conditional expectations of $z_{1,t}$ (that is, utility when a trade occurs) using the non-parametric kernel-based estimators described in section 3.2.2. These assume second moment timing ability (only) over the intra-daily, daily, weekly and bi-weekly trading sessions, with no transaction costs and a risk aversion (γ) level of eight.

Insert Figure 1 here

The plots clearly show that conditional expectations are non-normal. Perhaps most notable the distributions are truncated on the right-hand side, and also appear bi-modal when daily, weekly and bi-weekly trading sessions are considered. Thus the formula in Proposition 2 would seem to be

inappropriate because it relies on the normality assumption. Consequently, the subsequent analysis uses the expression in Proposition 1 to measure the expected utility associated with the ReDiT strategies.

3.3.2. *Estimating skill/Model performance*

A key determinant of performance is the ability (or otherwise) to predict each of the realised moments. To examine this skill we report the estimated fit ($\widehat{\mathcal{R}}_{k,n}^2$) associated with out-of-sample forecasts of each realised moment. These values are calculated under the assumption of intra-daily, daily, weekly and bi-weekly trading sessions. Statistical inference is achieved by performing a simple bootstrap procedure on the actual and forecasts to yield a bootstrap distribution for the $\widehat{\mathcal{R}}_{k,n}^2$ statistics.

Insert Table 2 here

The results are fairly clear cut regarding the performance of each set of expectations. First, naive expectations do not appear to be useful to investors in terms of forecasting realised moments. By contrast, adaptive and long-memory adaptive expectations offer substantial benefits. Of these it is the latter (which are based on the unrestricted FIMA model) that have the slight advantage. Second, only the second and, to a lesser extent, the fourth realised moments can be predicted. This result confirms the efficiency of energy futures markets in that none of the odd realised moments are predictable. Of the even realised moments, it is the second moment (realised variance) that is most predictable. For instance, over the daily trading session the $\widehat{\mathcal{R}}_{k,n}^2$ statistic associated with long-memory adaptive expectations equals 57.9%, which is significantly different from zero at the 1% level. Of the other realised moments, it is only the fourth realised moment that appears (significantly) predictable with long-memory adaptive expectations delivering an $\widehat{\mathcal{R}}_{k,n}^2$ statistic of 17.3% over the daily trading session.⁸ These results are consistent with the literature showing that returns are difficult to predict, while realised volatility (transformed second realised moment) is not. The results also show that there is a tendency for the predictability of the second and fourth moments to increase as the length of the trading session increases.

---

⁸Given the superior performance of the long-memory adaptive expectations, the subsequent analysis makes exclusive use of these expectations.

3.4. Strategy performance and risk aversion

The non-normal nature of the realised moments means that the result in Proposition 1 must be used to estimate expected performance (as opposed to using the formula in Proposition 2 based on the normality assumption). Using the procedure described in section 3.2.2 the expected realised session utilities associated with a variety of ReDiT strategies (and trading session lengths) are calculated. In particular, we consider ReDiT strategies that assume: i. no timing ability (all $\widehat{\mathcal{R}}_{k,n}^2$ values set to zero); ii. second moment timing (all $\widehat{\mathcal{R}}_{k,n}^2$ values set to zero except that associated with the second moment which is set to the value obtained using long-memory adaptive expectations); and iii. fourth moment timing (all $\widehat{\mathcal{R}}_{k,n}^2$ values set to zero except that associated with the fourth moment which is set to the value obtained using long-memory adaptive expectations). In addition, we consider the expected performance associated with the strategy characterised by always trading (henceforth the *always-trade* strategy). To provide a meaningful measure of expected performance it is converted to an annualised certainty equivalent return (CER) value. The results are provided in Table 3 for $\gamma = 0, 1, 2, 4, 8, 16$.

Insert Table 3 here

The results highlight a number of important features regarding the relative performance of ReDiT strategies. First, as risk aversion levels increase the performance of the always-trade strategy deteriorates rapidly. By contrast, the ReDiT strategies mitigate this performance downturn such that the CER values approach zero over this space – a result predicted in the previous section. Second, there is variation over the different ReDiT strategies. Of these strategies it is those based on timing the second realised moment that dominate. For instance, the ReDiT strategy based on second moment (volatility) timing employed during the daily trading session earns a CER of 1.10% when $\gamma = 4$. By contrast, all competing ReDiT strategies earn a zero CER.

The results also reveal variation over the trading session length. First, the always-trade and timing (unskilled) strategies deliver the same performance levels over all trading session lengths (hence the results are only reported for the intra-daily trading session). This is because realised moments over longer trading sessions are simply scaled versions of the realised moments observed over shorter trading sessions (e.g., the mean of the first moment observed over the weekly trading session is five times the mean of the first moment observed over the daily trading session). However,

the performance of the ReDiT strategies (with second or fourth moment timing ability) will vary of this space because the trading decision is different. In turn, this is due to variation in the parameter values of the prediction model (that is, the FIMA model and restricted versions thereof).

The results indicate that ReDiT strategy performance declines as the trading session lengthens. Thus despite the fact that there is improved predictability over lower frequencies (see Table 1), the realised moment variance per unit of time is sufficiently lower over longer trading sessions to yield a net deterioration in performance. This however does not rule out the use of longer trading sessions as these may be beneficial in terms of lower turnover levels and hence transaction costs. To examine this issue the relative CER values associated with the ReDiT strategy based on second moment timing (relative to the unskilled ReDiT strategy benchmark) are plotted against $\gamma \in [0, 8]$ for transaction costs levels of $\tau = 0, 0.0001, 0.0002, 0.0004, 0.0008$. These are provided in Figure 2.

Insert Figure 2 here

The results highlight the non-linear nature of relative performance with respect to risk aversion. In the absence of transaction costs, the ReDiT strategy achieves CER values up to 2% for γ values around three when shorter trading sessions are assumed. However, this value is highly sensitive to the risk aversion level, with performance falling sharply for alternative γ values below or above three. When transaction costs are introduced there are clear differences in the results over the different trading session lengths. For the shorter trading sessions there is a noticeable sharp decline in performance. Indeed, for transaction costs around eight basis points (that is, $\tau = 0.0008$), the benefits all but disappear for all risk aversion levels. By contrast, as the trading session length increases, the effect of transaction costs is greatly reduced. Indeed, when a bi-weekly trading session is assumed there is little difference between performance over the transaction costs. This finding reflects the lower turnover levels associated with longer trading sessions.

3.5. Dominant risk aversion regions

The previous results show that risk aversion is a key factor in determining whether the ReDiT strategy is useful. We can go a step further and calculate the range of risk aversion levels that deliver superior performance for the ReDiT strategy (based on second moment timing). These are referred to as *dominant risk aversion regions*. In particular, we present the range of risk aversion

values that deliver a relative CER value (that is, the CER value associated with the skilled ReDiT strategy minus the CER value associated with the unskilled ReDiT strategy) in excess of a pre-select target CER value. We consider annualised target values of 0.25%, 0.50%, 0.75% and 1.00%. Results associated with $\tau = 0, 0.0001, 0.0002, 0.0004, 0.0008$ are provided in Table 4.

Insert Table 4 here

In the absence of transaction costs we see that over the daily trading session the skilled ReDiT strategy attains the 0.25% target performance level for γ values between 2.0 and 7.3. As one increases the target level this range narrows such that a 1% target is consistent with γ values between 2.7 and 4.3. When transaction costs are introduced the dominant risk aversion region narrows. Indeed when $\tau = 0.0008$ there are no risk aversion levels consistent with achieving this target. As the trading session length is increased the dominant risk aversion region largely maintains its width in the presence of transaction costs.

The obvious question now is whether the dominant risk aversion regions observed in Table 4 include risk aversion levels associated with the typical investor. The upper value of the dominant risk aversion regions never exceeds 7.3 even with zero transaction costs. Two previous studies are useful in this regard. Specifically, Aït-Sahalia and Lo (2000) and Bliss and Panigirtzoglou (2004) provide empirical evidence that the risk aversion levels of the typical investor under power utility range from three to twelve. Thus the observed dominant risk aversion regions in Table 4 are realistic in the sense that they could potentially be relevant to the typical investor. However, the caveat here is that this consistency only holds for low transaction costs and target CER levels, or for long trading sessions.

3.6. Strategy performance and trading volume

There is growing evidence that moment prediction quality is dependent on market conditions; see, e.g., Taylor (2014a, 2014b) for evidence in the context of the first two moments. Following this line of enquiry we consider the performance of ReDiT strategies under different trading volume environments. During periods of low trading volume (low liquidity) we may expect transaction costs to be high; see Wang and Yau (2000) for empirical support in the context of futures markets. Consequently, we would expect ReDiT strategies to appear to perform well because transaction

cost levels are underestimated. By contrast, during periods of high trading volume (high liquidity) we would expect weaker performance as transaction costs are likely to be overestimated.

To investigate this hypothesis we divide the sample into five ascending trading volume quintiles.⁹ The performance of the ReDiT strategy based on second moment timing is then assessed for each trading volume quintile. Performance is measured as the average CER value (relative to the unskilled ReDiT strategy) obtained over the space $\gamma \in [3, 12]$ under the assumption of zero transaction costs.¹⁰ Results associated with use of intra-daily, daily, weekly and bi-weekly trading sessions are provided in Figure 3.

Insert Figure 3 here

The results are clear cut. There is a negative relationship between ReDiT strategy performance and trading volume. Low (high) trading volume environments are consistent with high (low) ReDiT strategy performance. Therefore, the superior performance of the ReDiT strategy during low trading volume (high transaction cost) environments is flattering and is likely to disappear if appropriate transaction cost levels are considered.

3.7. Required skill levels

The analysis assumes a particular investor skill level as dictated by the long-memory adaptive expectations. It is quite possible that investors may use an alternative expectation model that is associated with greater levels of skill. If this is the case then this naturally leads to the question of how much skill is required to achieve a particular level of performance *and* what are the risk aversion levels consistent with achieving this performance level? To answer these questions we consider hypothetical $\mathcal{R}_{k,n}^2$ levels of 0, 0.01, 0.1, and 0.5 (loosely speaking no, low, medium and high skill levels) and all integer risk aversion levels from zero to 300. The skill levels are assumed to be separately associated with first, second, third and fourth realised moment timers. For each of these risk aversion levels we calculate the skill level that must be exceeded to achieve an annualised target CER value of 0.5%.¹¹ Results associated with $\tau = 0, 0.0001, 0.0002, 0.0004, 0.0008$ are provided in

---

⁹Monthly trading volume data over the period 2007 to 2015 were obtained from Tickdata, Inc.

¹⁰This space most likely corresponds to the risk aversion levels of the typical investor (Aït-Sahalia and Lo, 2000, and Bliss and Panigirtzoglou, 2004)).

¹¹This target level coincides with the likely minimum target level as it corresponds to the lower end of riskfree rates observed over recent periods.

Table 5 (first realised moment timer), Table 6 (second realised moment timer) and Table 7 (third realised moment timer). Results associated with ReDiT strategies based on fourth realised moment timing are not presented as there is no combination of skill and risk aversion that will deliver the target CER value.

Insert Tables 5, 6 and 7 here

The results indicate that only very low skill levels are required when employing the ReDiT strategy based on first moment timing. Indeed, for risk aversion levels below twenty, skills levels above the no skill level are sufficient to yield the target CER. This result holds for all transaction costs. Only when risk aversion levels are extremely high is it impossible to achieve the target CER level (daily, weekly and bi-weekly trading sessions only).

For the ReDiT strategies based on higher moments, the required skill levels become higher. For volatility timers (that is, ReDiT strategies based on second moment timing) the required skill levels are high, but are below those associated with the long-memory adaptive expectations documented in Table 2. However, this result only holds in low transaction cost environments (intra-daily and daily trading sessions) or for a limited range of risk aversion levels (weekly and bi-weekly trading sessions). When ReDiT strategies based on higher moments are considered we note that the skill levels are very demanding or require a huge risk aversion level. For instance, for ReDiT strategies based on third moment timing applied during the daily trading session we require medium to high skill levels and a risk aversion level in excess of 171.

4. Conclusions

The framework developed in this paper enables us to decompose the drivers of timing strategy performance in the crude oil futures market. Perhaps the most obvious relates to the nature of the strategy itself. For ReDiT strategies based on the second realised moment (that is, volatility timers) then economically significant performance levels are obtained. For ReDiT strategies based on timing other moments we provide no evidence of significant performance. These results are consistent with the notion that the crude oil futures market is economically efficient. Other factors affect performance. Transaction costs are shown to have a large impact, though their effect can be mitigated through use of a weekly or bi-weekly trading session. Another consideration is the risk

aversion level of the investor. For instance, when using the recommended bi-weekly trading session, the benefits to volatility timing only accrue to investors with risk aversion levels between two and five.

The results in this paper should not deter those seeking speculative gains. Indeed, we provide a desiderata of skill levels for ReDiT strategies based on timing each moment under different transaction costs and risk aversion levels. For those who seek to time the first moment (that is, market timers) then the required skill levels are extremely small. This is because the utility function used by these investors places the greatest weight on this moment. However, these skill levels are the hardest to obtain as they violate market efficiency. As one moves to higher moments the required skill levels increase as the utility function places less weight on these moments. The choice of which moment to time will therefore ultimately depend on the investors' skill, risk aversion and transaction costs. The results in this paper provide guidance in this regard.

The investor framework is built on a number of assumptions. These are necessarily restrictive in order to yield compact formulae for expected performance under transaction costs. Future work could build on this framework by introducing alternative assumptions that involve a dynamic investment strategy, and/or different utility functions and/or different types of transaction costs (e.g. transaction costs that allow for market impact). The empirical section is also subject to assumptions. Perhaps the most noticeable is the choice of trading session lengths. One could consider even shorter trading sessions in order to assess the performance of low latency trading strategies. The framework developed in this paper is capable of allowing such choices.

Appendix

This appendix contains additional technical material including proofs of the results in section 2, and refined results based on additional assumptions.

A. Proofs

Proof of Proposition 1. It is useful to express utility in each state (and session) as follows

$$U_{k,t} = \begin{cases} z_{0,t}, & \text{if } x_{k,t-1} = 0 \text{ (no trade occurs),} \\ z_{1,t}, & \text{if } x_{k,t-1} = 1 \text{ (trade occurs),} \end{cases} \quad (\text{A.1})$$

where

$$z_{0,t} = \sum_{n=0}^N \beta_n \sum_{s=1}^S (-\kappa)^n = 0, \quad (\text{A.2a})$$

$$z_{1,t} = \sum_{n=0}^N \beta_n \sum_{m=0}^n \binom{n}{m} (\ln[1 - \tau])^{n-m} y_{m,t}. \quad (\text{A.2b})$$

Here $y_{m,t}$ is the m th realised moment associated with returns to the risky asset being traded. The expressions for $z_{0,t}$ and $z_{1,t}$ are obtained using the expression for realised session utility given by (6), with $x_{k,t-1}$ set to zero (to give $z_{0,t}$) and unity (to give $z_{1,t}$). If the conditional expectation of utility obtained from trading is greater than the conditional expectation of utility obtained when not trading then trading will take place. Otherwise, no trade will occur. It follows that the utility obtained during the t th trading session will be given by

$$U_{k,t} = 1_{\mathbb{E}[z_{0,t}|\mathcal{F}_{k,t-1}] \geq \mathbb{E}[z_{1,t}|\mathcal{F}_{k,t-1}]} z_{0,t} + 1_{\mathbb{E}[z_{1,t}|\mathcal{F}_{k,t-1}] > \mathbb{E}[z_{0,t}|\mathcal{F}_{k,t-1}]} z_{1,t}. \quad (\text{A.3})$$

As $z_{0,t} = 0$ we have

$$U_{k,t} = x_{k,t-1}^* z_{1,t}. \quad (\text{A.4})$$

where $x_{k,t-1}^*$ describes the optimal behaviour associated with the k th ReDiT strategy and is given by

$$x_{k,t-1}^* = 1_{\mathbb{E}[z_{1,t}|\mathcal{F}_{k,t-1}] > 0}. \quad (\text{A.5})$$

Here 1_c is an indicator function that equals unity if condition c holds, and zero otherwise. Thus optimal trading occurs according to (A.4), with the expression for $E[z_{1,t}|\mathcal{F}_{k,t-1}]$ obtained by taking conditional expectations of (A.2b).

Using the expression for $z_{1,t}$ in (A.2b) and decomposing via (9) we obtain

$$\begin{aligned} U_{k,t} &= x_{k,t-1}^* \sum_{n=0}^N \beta_n \sum_{m=0}^n \binom{n}{m} (\ln[1-\tau])^{n-m} y_{m,t}, \\ &= x_{k,t-1}^* \sum_{n=0}^N \beta_n \sum_{m=0}^n \binom{n}{m} (\ln[1-\tau])^{n-m} (E[y_{m,t}|\mathcal{F}_{k,t-1}] + \epsilon_{k,m,t}). \end{aligned} \quad (\text{A.6})$$

Taking expectations and noting that expectations involving $\epsilon_{k,t}$ equal zero gives

$$E[U_{k,t}] = E[x_{k,t-1}^* E[z_{1,t}|\mathcal{F}_{k,t-1}]]. \quad (\text{A.7})$$

Using the result in (A.5) leads to

$$\begin{aligned} E[U_{k,t}] &= E[1_{E[z_{1,t}|\mathcal{F}_{k,t-1}]>0} E[z_{1,t}|\mathcal{F}_{k,t-1}]], \\ &= \Pr[E[z_{1,t}|\mathcal{F}_{k,t-1}] > 0] \times E[E[z_{1,t}|\mathcal{F}_{k,t-1}] | E[z_{1,t}|\mathcal{F}_{k,t-1}] > 0], \\ &= (1 - G_{E[z_{1,t}|\mathcal{F}_{k,t-1}]}[0]) \times \frac{\int_{w=0}^{\infty} w g_{E[z_{1,t}|\mathcal{F}_{k,t-1}]}[w] dw}{(1 - G_{E[z_{1,t}|\mathcal{F}_{k,t-1}]}[0])}, \end{aligned} \quad (\text{A.8})$$

where $G_{E[z_{1,t}|\mathcal{F}_{k,t-1}]}[.]$ is the cumulative distribution function associated with $E[z_{1,t}|\mathcal{F}_{k,t-1}]$. Simplifying leads to the result in Proposition 1. $\square$

B. Further results

To enable refinement of the result in Proposition 1 a number of additional assumptions are required. These are provided below.

Assumption 9. Individual realised moments and forecast errors have (independent) normal distributions: $y_{n,t} \sim \mathcal{N}(\mu_n, \sigma_n^2)$ and $\epsilon_{k,n,t} \sim \mathcal{N}(0, \sigma_{k,n}^2(\epsilon))$.

Remark. It follows that $E[y_{n,t}|\mathcal{F}_{k,t-1}] \sim \mathcal{N}(\mu_n, \sigma_n^2 - \sigma_{k,n}^2(\epsilon))$. Here μ_n and σ_n^2 are exogenously determined, while $\sigma_{k,n}^2(\epsilon)$ is under the control of the investor. Furthermore, the coefficient of

determination associated with each moment is given by

$$\mathcal{R}_{k,n}^2 = 1 - \frac{\sigma_{k,n}^2(\epsilon)}{\sigma_n^2} = \frac{\sigma_n^2 - \sigma_{k,n}^2(\epsilon)}{\sigma_n^2}. \quad (\text{B.1})$$

This expression enables us to rewrite the variance of $\mathbb{E}[y_{n,t}|\mathcal{F}_{k,t-1}]$ such that $\mathbb{E}[y_{n,t}|\mathcal{F}_{k,t-1}] \sim \mathcal{N}(\mu_n, \sigma_n^2 \mathcal{R}_{k,n}^2)$.

Assumption 10. Individual realised moments are independent of each other such that $\text{cov}[y_{m,t}, y_{n,t}] = 0 \forall m \neq n$.

Proposition 2. *Under Assumptions 1 to 10, the unconditional expectation of the realised session utility to the k th ReDiT strategy can be decomposed as follows:*

$$\mathbb{E}[U_{k,t}] = \mu \Phi[\alpha_k] + \sqrt{\Omega_k} \phi[\alpha_k], \quad (\text{B.2})$$

where $\alpha_k = \mu/\sqrt{\Omega_k}$, with

$$\mu = \sum_{n=0}^N \beta_n \sum_{m=0}^n \binom{n}{m} (\ln[1 - \tau])^{n-m} \mu_m, \quad (\text{B.3a})$$

$$\Omega_k = \sum_{n=0}^N \beta_n^2 \sum_{m=0}^n \binom{n}{m} \binom{n}{m} (\ln[1 - \tau])^{2(n-m)} \sigma_m^2 \mathcal{R}_{k,m}^2. \quad (\text{B.3b})$$

Here $\phi[\cdot]$ and $\Phi[\cdot]$ are the standard normal probability density and cumulative distribution functions, respectively. The measure μ represents the total reward to trading, and Ω_k is the total investor skill.

Proof. As $\mathbb{E}[z_{1,t}|\mathcal{F}_{k,t-1}] \sim \mathcal{N}(\mu, \Omega_k)$ then standard results associated with truncated normal variables lead directly to the result in Proposition 2. □

Corollary 1. *The relationship between ReDiT strategy performance and skill can be summarised*

as follows:

$$\lim_{\Omega_k \rightarrow 0} \mathbb{E}[U_{k,t}] = \max[0, \mu] \geq 0, \quad (\text{B.4a})$$

$$\frac{\partial \mathbb{E}[U_{k,t}]}{\partial \Omega_k} = \frac{1}{2\sqrt{\Omega_k}} \phi[\alpha_k] \geq 0, \quad (\text{B.4b})$$

$$\frac{\partial^2 \mathbb{E}[U_{k,t}]}{\partial \Omega_k^2} = \frac{\mu^2 - \Omega_k}{4\sqrt{\Omega_k^5}} \phi[\alpha_k] \gtrless 0 \quad \text{if } \Omega_k \lesseqgtr \mu^2. \quad (\text{B.4c})$$

Here the first equation describes the absence of skill case, while the second and third equation describe how additional skill affects performance.

Proof. The results follow from Proposition 2. □

Remark. The first equation shows that in the absence of skill there is still potential benefit to using ReDiT strategies over always trading. This is because when the rewards to trading are negative (that is, $\mu < 0$), then the ReDiT strategy selects the no trade option. By contrast, when there are positive rewards to trading (that is, $\mu > 0$) then the ReDiT strategy selects to trade option. The second equation shows that additional skill is always rewarded. Finally, the third equation shows that if $\mu \neq 0$ then the marginal benefits of additional skill are positive when the initial skill levels are low. Otherwise, the marginal benefits of additional skill are negative.

References

- Aït-Sahalia, Y., Lo, A., 2000. Nonparametric risk management and implied risk aversion. *Journal of Econometrics* 94, 9–51.
- Bliss, R., Panigirtzoglou, N., 2004. Option-implied risk aversion estimates. *Journal of Finance* 59, 407–446.
- Chiriac, R., Voev, V., 2011. Modelling and forecasting multivariate realised volatility. *Journal of Applied Econometrics* 26, 922–947.
- Fleming, J., Kirby, C., Ostdiek, B., 2001. The economic value of volatility timing. *Journal of Finance* 56, 329–352.
- Fleming, J., Kirby, C., Ostdiek, B., 2003. The economic value of volatility timing using realized volatility. *Journal of Financial Economics* 67, 473–509.
- Garlappi, L., Skoulakis, G., 2011. Taylor series approximations to expected utility and optimal portfolio choice. *Mathematics and Financial Economics* 5, 121–156.
- Jondeau, E., Rockinger, M., 2012. On the importance of time variability in higher moments for asset allocation. *Journal of Financial Econometrics* 10, 84–123.
- Jones, C., 2013, What do we know about high-frequency trading? Columbia Business School Research Paper No. 13-11.
- Kandel, S., Stambaugh, R., 1996. On the predictability of stock returns: An asset allocation perspective. *Journal of Finance* 51, 385–424.
- Kang, S., McIver, R., Yoon, S., 2017. Dynamic spillover effects among crude oil, precious metal, and agricultural commodity futures markets. *Energy Economics* 62, 19–32.
- Kristoufek, L., Vosvrda, M., 2014. Commodity futures and market efficiency. *Energy Economics* 42, 50–57.
- Levy, H., Markowitz, H., 1979. Approximating expected utility as a function of mean and variance. *American Economic Review* 69, 308–317.

- Liu, L., Patton, A., Sheppard, K., 2015. Does anything beat 5-minute RV? A comparison of realized measures across multiple asset classes. *Journal of Econometrics* 187, 293–311.
- Liu, X., An, H., Wang, L., 2015. Performance of generated moving average strategies in natural gas futures prices at different time scales. *Journal of Natural Gas Science and Engineering* 24, 337–345.
- Liu, X., An, H., Wang, L., Jia, X., 2017a. An integrated approach to optimize moving average rules in the EUA futures market based on particle swarm optimization and genetic algorithms. *Applied Energy* 185, 1778–1787.
- Liu, X., An, H., Wang, L., Guan, Q., 2017b. Quantified moving average strategy of crude oil futures market based on fuzzy logic rules and genetic algorithms. *Physica A: Statistical Mechanics and its Applications* 482, 444–457.
- Lubnau, T., Todorova, N., 2015. Trading on mean-reversion in energy futures markets. *Energy Economics* 51, 312–319.
- Marquering, W., Verbeek, M., 2004. The economic value of predicting stock index returns and volatility. *Journal of Financial and Quantitative Analysis* 39, 407–429.
- Merton, R., 1981. On market timing and investor performance. I. An equilibrium theory of value for market forecasts. *Journal of Business* 54, 363–406.
- Proietti, T., 2016. Component-wise representations of long-memory models and volatility prediction. *Journal of Financial Econometrics* 14, 668–692.
- Silverman, B., 1986. *Density Estimation for Statical and Data Analysis*, Chapman and Hall, London.
- Taylor, N., 2014a. The economic value of volatility forecasts: A conditional approach. *Journal of Financial Econometrics* 12, 433–478.
- Taylor, N., 2014b. The rise and fall of technical trading rule success. *Journal of Banking and Finance* 40, 286–302.

- Wang, G., Yau, J., 2000. Trading volume, bidask spread, and price volatility in futures markets. *Journal of Futures Markets* 20, 943–970.
- Wang, L., An, H., Liu, X., Huang, X., 2016a. Selecting dynamic moving average trading rules in the crude oil futures market using a genetic approach. *Applied Energy* 162, 1608–1618.
- Wang, T., Yang, J., 2010. Nonlinearity and intraday efficiency tests on energy futures markets. *Energy Economics* 32, 496–503.
- Wang, Y., Liu, L., Ma, F., Wu., C., 2016b. What the investors need to know about forecasting oil futures return volatility. *Energy Economics* 57, 128–139.
- Welch, I., Goyal, A., 2008. A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies* 21, 1455–1508.
- West, K., Edison, H., Cho, D., 1993. A utility-based comparison of some models of exchange rate volatility. *Journal of International Economics* 35, 23–45.

Table 1 – Summary statistics

		Moment					
Measure		1st	2nd	3rd	4th	5th	6th
Panel A: Intra-daily trading session							
Mean:	μ_n	1.65×10^{-4}	1.50×10^{-4}	-6.00×10^{-8}	7.77×10^{-9}	-1.88×10^{-11}	2.21×10^{-12}
Standard deviation:	σ_n	1.20×10^{-2}	1.73×10^{-4}	1.48×10^{-6}	4.44×10^{-8}	1.15×10^{-9}	3.96×10^{-11}
Skewness:		-8.38×10^{-3}	4.56×10^0	-8.14×10^{-3}	2.19×10^1	2.28×10^1	5.41×10^1
Kurtosis (excess):		4.09×10^0	3.38×10^1	3.37×10^2	7.36×10^2	2.65×10^3	3.84×10^3
Panel B: Daily trading session							
Mean:	μ_n	3.31×10^{-4}	3.00×10^{-4}	-1.20×10^{-7}	1.55×10^{-8}	-3.76×10^{-11}	4.42×10^{-12}
Standard deviation:	σ_n	1.71×10^{-2}	3.15×10^{-4}	2.10×10^{-6}	6.94×10^{-8}	1.64×10^{-9}	5.74×10^{-11}
Skewness:		-1.25×10^{-2}	4.10×10^0	1.24×10^0	1.47×10^1	1.68×10^1	3.65×10^1
Kurtosis (excess):		3.60×10^0	2.59×10^1	1.98×10^2	3.04×10^2	1.30×10^3	1.76×10^3
Panel C: Weekly trading session							
Mean:	μ_n	1.65×10^{-3}	1.50×10^{-3}	-6.00×10^{-7}	7.77×10^{-8}	-1.88×10^{-10}	2.21×10^{-11}
Standard deviation:	σ_n	3.87×10^{-2}	1.34×10^{-3}	4.73×10^{-6}	2.15×10^{-7}	3.52×10^{-9}	1.37×10^{-10}
Skewness:		-1.97×10^{-1}	3.41×10^0	-5.51×10^{-1}	6.80×10^0	8.35×10^0	1.47×10^1
Kurtosis (excess):		4.22×10^0	1.65×10^1	3.85×10^1	5.59×10^1	2.85×10^2	2.88×10^2
Panel D: Bi-weekly trading session							
Mean:	μ_n	3.31×10^{-3}	3.00×10^{-3}	-1.20×10^{-6}	1.55×10^{-7}	-3.76×10^{-10}	4.42×10^{-11}
Standard deviation:	σ_n	5.53×10^{-2}	2.56×10^{-3}	6.73×10^{-6}	3.86×10^{-7}	4.97×10^{-9}	2.14×10^{-10}
Skewness:		-5.75×10^{-1}	3.29×10^0	-1.23×10^0	6.06×10^0	5.19×10^0	9.76×10^0
Kurtosis (excess):		2.60×10^0	1.55×10^1	2.23×10^1	4.35×10^1	1.33×10^2	1.20×10^2

Notes: This table contains the mean, standard deviation, skewness and kurtosis (excess) of the first six realised moments. The sample period used is January 1, 1992, to December 31, 2015.

Table 2 – Investor skill with respect to individual realised moments

Forecasting Model	Moment					
	1st	2nd	3rd	4th	5th	6th
Panel A: Intra-daily trading session						
M0: Naive	−0.001	−0.004	0.001	−0.009	0.000	−0.003
M1: Adaptive	−0.001	0.493**	−0.001	0.089**	−0.001	0.004
M2: Long-memory Adaptive	−0.002	0.504**	−0.005	0.088**	−0.002	0.004
Panel B: Daily trading session						
M0: Naive	−0.001	−0.005	0.001	−0.015	−0.001	−0.005
M1: Adaptive	−0.001	0.568**	−0.002	0.144**	−0.002	0.010
M2: Long-memory Adaptive	−0.005	0.579**	−0.006	0.173**	−0.003	0.016
Panel C: Weekly trading session						
M0: Naive	−0.006	−0.011	0.007	−0.041	−0.004	−0.022
M1: Adaptive	−0.004	0.695**	−0.001	0.427**	−0.008	0.097
M2: Long-memory Adaptive	−0.019	0.690**	−0.043	0.444**	−0.016	0.107
Panel D: Bi-weekly trading session						
M0: Naive	−0.011	−0.016	0.014	−0.055	−0.009	−0.039
M1: Adaptive	−0.010	0.710**	0.003	0.530**	−0.014	0.075
M2: Long-memory Adaptive	−0.017	0.705**	−0.039	0.519**	−0.033	0.060

Notes: This table contains the fit (coefficient of determination, $\mathcal{R}_{k,n}^2$) associated with out-of-sample forecasts of each realised moment. Forecasts are based on models with parameters estimated using a rolling window of past observations. Cases where the basic bootstrap 95% (99%) confidence interval does not zero are indicated by * (**).

Table 3 – Strategy performance and risk aversion

Strategy	Risk Aversion (γ)					
	0	1	2	4	8	16
Panel A: Intra-daily trading session						
Always-trade	12.013	8.267	4.519	-2.979	-17.990	-48.093
ReDiT (unskilled)	12.013	8.267	4.519	0.000	0.000	0.000
ReDiT (second moment timing)	12.023	8.267	4.792	1.189	0.193	0.014
ReDiT (fourth moment timing)	12.013	8.267	4.519	0.000	0.000	0.000
Panel B: Daily trading session						
ReDiT (second moment timing)	12.028	8.267	4.765	1.099	0.112	0.000
ReDiT (fourth moment timing)	12.013	8.267	4.519	0.000	0.000	0.000
Panel C: Weekly trading session						
ReDiT (second moment timing)	12.062	8.267	4.670	0.838	0.014	0.000
ReDiT (fourth moment timing)	12.013	8.267	4.519	0.000	0.000	0.000
Panel D: Bi-weekly trading session						
ReDiT (second moment timing)	12.069	8.267	4.629	0.719	0.000	0.000
ReDiT (fourth moment timing)	12.013	8.267	4.519	0.000	0.000	0.000

Notes: This table contains the annualised CER values associated with various trading strategies. The timing strategies differentiate themselves in terms of the realised moments on which the trade decision is made. Forecasts are based on the long-memory adaptive model with parameters estimated using a rolling window of past observations.

Table 4 – Risk aversion levels for dominant timing strategies

Target CER	Transaction Cost (τ)									
	0×10^{-4}		1×10^{-4}		2×10^{-4}		4×10^{-4}		8×10^{-4}	
	L	U	L	U	L	U	L	U	L	U
Panel A: Intra-daily trading session										
$\geq 0.25\%$	2.0	7.3	1.6	2.6	0.4	0.6	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$
$\geq 0.50\%$	2.3	5.6	1.8	2.1	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$
$\geq 0.75\%$	2.5	4.8	1.9	1.9	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$
$\geq 1.00\%$	2.7	4.3	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$
Panel B: Daily trading session										
$\geq 0.25\%$	2.0	6.5	1.8	4.4	1.6	2.5	0.4	0.6	$\emptyset$	$\emptyset$
$\geq 0.50\%$	2.4	5.3	2.1	3.5	1.8	2.1	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$
$\geq 0.75\%$	2.6	4.6	2.3	3.1	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$
$\geq 1.00\%$	2.8	4.1	2.4	2.8	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$
Panel C: Weekly trading session										
$\geq 0.25\%$	2.1	5.4	2.1	5.1	2.1	4.7	2.0	4.1	1.8	2.9
$\geq 0.50\%$	2.5	4.6	2.5	4.3	2.4	4.0	2.3	3.5	2.0	2.4
$\geq 0.75\%$	2.7	4.1	2.7	3.8	2.6	3.6	2.5	3.1	2.1	2.1
$\geq 1.00\%$	2.9	3.7	2.8	3.5	2.8	3.2	2.6	2.8	$\emptyset$	$\emptyset$
Panel D: Bi-weekly trading session										
$\geq 0.25\%$	2.2	5.0	2.2	4.9	2.2	4.7	2.1	4.4	2.1	3.8
$\geq 0.50\%$	2.6	4.3	2.6	4.2	2.5	4.1	2.5	3.8	2.3	3.3
$\geq 0.75\%$	2.8	3.9	2.8	3.8	2.7	3.7	2.7	3.4	2.5	3.0
$\geq 1.00\%$	3.0	3.6	2.9	3.5	2.9	3.4	2.8	3.1	2.6	2.7

Notes: This table contains the range (lower, L, and upper, U, values) of risk aversion (γ) levels in which the ReDiT strategy (based on the second realised moment) is dominant. This is calculated for various cutoffs values (in annualised percentage terms) and transaction cost levels. A strategy is dominant if it beats the always-trade and do-nothing strategies by at least the cutoff value. Cases where dominance is not achieved at any risk aversion level are denoted by $\emptyset$.

Table 5 – Required skill levels (first realised moment timer)

Risk Aversion (γ)	Transaction Cost (τ)				
	0×10^{-4}	1×10^{-4}	2×10^{-4}	4×10^{-4}	8×10^{-4}
Panel A: Intra-daily trading session					
[0, 40]	> 0	> 0	> 0	> 0	> 0
[41, 45]	> 0	> 0	> 0	> 0	> 0.01
[46, 47]	> 0	> 0	> 0	> 0.01	> 0.01
[48, 48]	> 0	> 0	> 0.01	> 0.01	> 0.01
[49, 49]	> 0	> 0.01	> 0.01	> 0.01	> 0.01
[50, 182]	> 0.01	> 0.01	> 0.01	> 0.01	> 0.01
[183, 300]	> 0.1	> 0.1	> 0.1	> 0.1	> 0.1
Panel B: Daily trading session					
[0, 21]	> 0	> 0	> 0	> 0	> 0
[22, 24]	> 0	> 0	> 0	> 0	> 0.01
[25, 25]	> 0	> 0	> 0	> 0.01	> 0.01
[26, 27]	> 0	> 0	> 0.01	> 0.01	> 0.01
[27, 95]	> 0.01	> 0.01	> 0.01	> 0.01	> 0.01
[96, 96]	> 0.01	> 0.01	> 0.01	> 0.01	> 0.1
[97, 97]	> 0.01	> 0.01	> 0.01	> 0.1	> 0.1
[98, 98]	> 0.01	> 0.01	> 0.1	> 0.1	> 0.1
[99, 203]	> 0.1	> 0.1	> 0.1	> 0.1	> 0.1
[204, 204]	> 0.5	> 0.1	> 0.1	> 0.1	> 0.1
[205, 205]	> 0.5	> 0.5	> 0.5	> 0.5	> 0.1
[206, 255]	> 0.5	> 0.5	> 0.5	> 0.5	> 0.5
[256, 256]	impos.	> 0.5	> 0.5	> 0.5	> 0.5
[257, 257]	impos.	impos.	> 0.5	> 0.5	> 0.5
[258, 258]	impos.	impos.	impos.	> 0.5	> 0.5
[259, 261]	impos.	impos.	impos.	impos.	> 0.5
[262, 300]	impos.	impos.	impos.	impos.	impos.
Panel C: Weekly trading session					
[0, 8]	> 0	> 0	> 0	> 0	> 0
[9, 9]	> 0	> 0	> 0	> 0	> 0.01
[10, 32]	> 0.01	> 0.01	> 0.01	> 0.01	> 0.01
[33, 33]	> 0.01	> 0.01	> 0.01	> 0.1	> 0.1
[34, 84]	> 0.1	> 0.1	> 0.1	> 0.1	> 0.1
[85, 85]	> 0.5	> 0.5	> 0.1	> 0.1	> 0.1
[86, 124]	> 0.5	> 0.5	> 0.5	> 0.5	> 0.5
[125, 125]	impos.	impos.	impos.	impos.	> 0.5
[126, 300]	impos.	impos.	impos.	impos.	impos.
Panel D: Bi-weekly trading session					
[0, 6]	> 0	> 0	> 0	> 0	> 0
[7, 7]	> 0	> 0	> 0	> 0.01	> 0.01
[8, 20]	> 0.01	> 0.01	> 0.01	> 0.01	> 0.01
[21, 21]	> 0.01	> 0.01	> 0.01	> 0.1	> 0.1
[22, 50]	> 0.1	> 0.1	> 0.1	> 0.1	> 0.1
[51, 74]	> 0.5	> 0.5	> 0.5	> 0.5	> 0.5
[75, 300]	impos.	impos.	impos.	impos.	impos.

Notes: This table contains the minimum $\mathcal{R}_{k,n}^2$ values (in the forecasting test equation) such that the ReDiT strategy (based on first realised moment timing) is dominant. A strategy is dominant if it beats the always-trade and do-nothing strategies by the target CER value of 0.5%. Cases where an $\mathcal{R}_{k,n}^2$ equal to unity is insufficient for dominance is referred to as ‘impossible’ (impos.).

Table 6 – Required skill levels (second realised moment timer)

Risk Aversion (γ)	Transaction Cost (τ)				
	0×10^{-4}	1×10^{-4}	2×10^{-4}	4×10^{-4}	8×10^{-4}
Panel A: Intra-daily trading session					
[0, 1]	impos.	impos.	impos.	impos.	impos.
[2, 2]	> 0.5	> 0.1	impos.	impos.	impos.
[3, 3]	> 0.01	> 0.5	impos.	impos.	impos.
[4, 5]	> 0.1	> 0.5	impos.	impos.	impos.
[6, 21]	> 0.5	> 0.5	impos.	impos.	impos.
[22, 41]	> 0.5	> 0.5	> 0.5	impos.	impos.
[42, 84]	> 0.5	> 0.5	> 0.5	> 0.5	impos.
[85, 134]	> 0.5	> 0.5	> 0.5	> 0.5	> 0.5
[135, 157]	> 0.5	> 0.5	> 0.5	> 0.5	impos.
[158, 164]	> 0.5	> 0.5	> 0.5	impos.	impos.
[165, 167]	> 0.5	> 0.5	impos.	impos.	impos.
[168, 170]	> 0.5	impos.	impos.	impos.	impos.
[171, 300]	impos.	impos.	impos.	impos.	impos.
Panel B: Daily trading session					
[0, 1]	impos.	impos.	impos.	impos.	impos.
[2, 2]	impos.	> 0.5	> 0.1	impos.	impos.
[3, 3]	> 0.01	> 0.1	impos.	impos.	impos.
[4, 5]	> 0.1	> 0.5	impos.	impos.	impos.
[6, 9]	> 0.5	> 0.5	impos.	impos.	impos.
[10, 38]	> 0.5	impos.	impos.	impos.	impos.
[39, 81]	> 0.5	> 0.5	impos.	impos.	impos.
[82, 93]	> 0.5	impos.	impos.	impos.	impos.
[94, 300]	impos.	impos.	impos.	impos.	impos.
Panel C: Weekly trading session					
[0, 1]	impos.	impos.	impos.	impos.	impos.
[2, 2]	impos.	impos.	impos.	impos.	> 0.1
[3, 3]	> 0.1	> 0.01	> 0.1	> 0.1	> 0.5
[4, 4]	> 0.1	> 0.1	> 0.5	> 0.5	impos.
[5, 5]	> 0.5	> 0.5	> 0.5	> 0.5	impos.
[6, 6]	> 0.5	> 0.5	> 0.5	impos.	impos.
[7, 7]	> 0.5	impos.	impos.	impos.	impos.
[8, 300]	impos.	impos.	impos.	impos.	impos.
Panel D: Bi-weekly trading session					
[0, 2]	impos.	impos.	impos.	impos.	impos.
[3, 3]	> 0.1	> 0.1	> 0.01	> 0.1	> 0.1
[4, 4]	> 0.1	> 0.1	> 0.5	> 0.5	> 0.5
[5, 5]	> 0.5	> 0.5	> 0.5	> 0.5	impos.
[6, 6]	> 0.5	> 0.5	impos.	impos.	impos.
[7, 300]	impos.	impos.	impos.	impos.	impos.

Notes: This table contains the minimum $\mathcal{R}_{k,n}^2$ values (in the forecasting test equation) such that the ReDiT strategy (based on second realised moment timing) is dominant. A strategy is dominant if it beats the always-trade and do-nothing strategies by the target CER value of 0.5%. Cases where an $\mathcal{R}_{k,n}^2$ equal to unity is insufficient for dominance is referred to as ‘impossible’ (impos.).

Table 7 – Required skill levels (third realised moment timer)

Risk Aversion (γ)	Transaction Cost (τ)				
	0×10^{-4}	1×10^{-4}	2×10^{-4}	4×10^{-4}	8×10^{-4}
Panel A: Intra-daily trading session					
[0, 84]	impos.	impos.	impos.	impos.	impos.
[85, 85]	> 0.5	> 0.5	> 0.5	impos.	impos.
[86, 86]	> 0.5	> 0.5	> 0.5	> 0.5	impos.
[87, 108]	> 0.5	> 0.5	> 0.5	> 0.5	> 0.5
[109, 109]	> 0.1	> 0.1	> 0.1	> 0.1	> 0.5
[110, 202]	> 0.1	> 0.1	> 0.1	> 0.1	> 0.1
[203, 206]	> 0.1	> 0.1	> 0.1	> 0.1	> 0.01
[207, 208]	> 0.1	> 0.1	> 0.1	> 0.01	> 0.01
[209, 209]	> 0.1	> 0.1	> 0.01	> 0.01	> 0.01
[210, 210]	> 0.1	> 0.01	> 0.01	> 0.01	> 0.01
[211, 300]	> 0.01	> 0.01	> 0.01	> 0.01	> 0.01
Panel B: Daily trading session					
[0, 170]	impos.	impos.	impos.	impos.	impos.
[171, 171]	impos.	impos.	impos.	impos.	> 0.5
[172, 172]	impos.	impos.	impos.	> 0.5	> 0.5
[173, 260]	> 0.5	> 0.5	> 0.5	> 0.5	> 0.5
[261, 275]	> 0.5	> 0.5	> 0.5	> 0.5	> 0.1
[276, 287]	> 0.5	> 0.5	> 0.5	> 0.1	> 0.1
[288, 297]	> 0.5	> 0.5	> 0.1	> 0.1	> 0.1
[298, 300]	> 0.5	> 0.1	> 0.1	> 0.1	> 0.1
Panel C: Weekly trading session					
[0, 282]	impos.	impos.	impos.	impos.	impos.
[283, 300]	impos.	impos.	impos.	impos.	> 0.5
Panel D: Bi-weekly trading session					
[0, 300]	impos.	impos.	impos.	impos.	impos.

Notes: This table contains the minimum $\mathcal{R}_{k,n}^2$ values (in the forecasting test equation) such that the ReDiT strategy ((based on third realised moment timing)) is dominant. A strategy is dominant if it beats the always-trade and do-nothing strategies by the target CER value of 0.5%. Cases where an $\mathcal{R}_{k,n}^2$ equal to unity is insufficient for dominance is referred to as ‘impossible’ (impos.).

Figure 1 – The distribution of conditional expectations

This figure provides kernel-based density plots of the conditional expectations of $z_{1,t}$ (that is, utility when a trade occurs) associated with second moment timing. Zero transaction costs and a risk aversion level (γ) of eight are assumed.

Figure 2 – Performance and risk aversion

This figure provides plots of ReDiT strategy (based on second moment timing) performance given by the annualised percentage CER with respect to the unskilled ReDiT strategy benchmark against risk aversion. Transaction cost levels from zero to 8×10^{-4} are assumed.

Figure 3 – Performance and trading volume

This figure provides plots of ReDiT strategy (based on second moment timing) performance given by the annualised percentage CER with respect to the unskilled ReDiT strategy benchmark. The CER value represents the average relative CER value obtained over the space $\gamma \in [3, 12]$. These values are plotted for each trading volume quintile. The dashed line represents the average CER value obtained over all quintiles. Zero transaction cost levels are assumed.