Taylor, N. J. (2016). Roll strategy efficiency in commodity futures markets. *Journal of Commodity Markets*, 1(1), 14-34.
<https://doi.org/10.1016/j.jcomm.2015.12.001>

Peer reviewed version

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

Link to published version (if available):
[10.1016/j.jcomm.2015.12.001](https://doi.org/10.1016/j.jcomm.2015.12.001)

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

This is the author accepted manuscript (AAM). The final published version (version of record) is available online via Elsevier at <http://www.sciencedirect.com/science/article/pii/S2405851315300143>. 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/>

Roll strategy efficiency in commodity futures markets[☆]

Nick Taylor^a

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

Abstract

Issues pertaining to the investor decision to sell a security and buy another (of the same type and with the same terms) with a longer period until the expiration date (the roll forward decision) are examined. In particular, a framework is developed in which it is possible to test the trade execution quality efficiency of a roll strategy against a mean-variance optimal roll strategy characterized by multiple-day roll. Applying this framework to five leading US grain futures markets (corn, wheat, soybean, soybean meal and soybean oil) demonstrates that commonly used single-day and multiple-day roll strategies (including the Goldman roll strategy) exhibit considerable inefficiencies. These are consistent over the markets and over the time of the day in which trading occurs, and vary with execution quality risk-aversion in a predictable way. A practical multiple-day roll strategy is proposed that reduces these inefficiencies.

Keywords: Roll strategy, execution risk, Bayesian inference, Goldman roll.

JEL: C11, G11, G13.

[☆]I thank participants in the European Financial Management Association Conference at Nyenrode University and in the departmental finance seminar at the University of Strathclyde for helpful comments.

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

1. Introduction

A large number of investment strategies involve the decision of when to sell a security and buy another (of the same type and with the same terms) with a longer period until the expiration date (the roll forward decision). For instance, a portfolio manager with responsibility for a diversified portfolio of assets that includes a commodity market component will undoubtedly have positions in commodity derivative contracts. Consequently, the manager must decide when to roll forward the derivative contracts which are close to maturity. Likewise, a stack hedger with a risk exposure to a particular commodity market that extends beyond the maturity of available derivative contracts faces a similar decision. A key factor in such decisions is expected trade execution quality. The approach taken in the current paper considers this decision such that this component is optimized.

There are a number of conventional approaches to the roll forward decision, all of which exploit the *periodic* nature of expected execution quality in the period leading up to contract maturity (cf. the vast majority of trade scheduling environments). These can be categorized as either single-day or multiple-day strategies. The former assumes that the securities roll forward on a single day that is a fixed number of days prior to the maturity date. By contrast, multiple-day strategies involve spreading the roll forward over a finite number of days prior to the maturity date. For instance, the commonly used *Goldman roll strategy* involves a commodity futures roll (with uniform weights) between the fifth and ninth business days of the month preceding the expiration month. The likely virtue of such a strategy in comparison to the single-day strategy is that the user has a lower risk of being exposed to poor execution quality. It is this conjecture that is examined in the current paper.

The advent of electronic trading platforms has led to an increasing number of financial institutions employing algorithmic trading systems. A key component of these systems is trade scheduling, defined according to a trade target, that is, the number of asset units (shares or contracts) to be bought or sold during a pre-specified finite time horizon. Specifically, trade scheduling involves specifying the rate at which these asset units are traded over this period (referred to as the trade list) in order to optimize execution quality. Within the context of derivative markets, a simple-to-implement trade scheduling procedure is proposed that solves the problem of when traders in such markets should switch from contracts that are close to maturity to deferred contracts (hence-

forth referred to as the roll strategy). This strategy applies irrespective of whether the user has speculative, hedging, or arbitrage motives.

Trade scheduling is important as it is a significant determinant of the overall success of a trading strategy. The academic literature has recognized this importance. The vast majority of studies propose strategies that optimize the tradeoff between pricing impact and timing risk; see Almgren and Chriss (2001), He and Mamaysky (2005), Engle and Ferstenberg (2007), Schied and Schoeneborn (2009), Forsyth et al. (2012), and Tse et al. (2013) for a representative sample.¹ This literature is complemented by proposing a framework that is designed specifically to examine the quality of the roll decision faced by participants in markets in which there is periodic variation in execution quality. Our framework has some overlap with this literature, but fundamentally differs in terms of the tradeoff undertaken.

The trade scheduling literature typically adopts the following framework. A trader wishes to sell a fixed number of asset units over a finite horizon. Execution quality is typically measured by comparing the total revenue generated by selling at the arrival price (the price observed when the trade instruction is received) and the total revenue generated by selling these units over the horizon, with the difference referred to as the implementation shortfall (Perold 1988). Within this context, price impact is generated by assuming that the asset is subject to trading costs that increase disproportionately with the trading rate (for instance, most studies adopt a quadratic cost model). This provides the incentive to avoid fast liquidation. By contrast, timing risk represents the risk of trading at prices away from the arrival price (induced by assuming that prices evolve in a stochastic fashion). These two effects are commonly balanced such that execution quality is optimized within a mean-variance framework in which the expected implementation shortfall is minimized subject to a pre-specified implementation shortfall variance. Moreover, these moments are determined from the perspective of a single trader prior to the trading taking place.²

Within our setup a trader holds a position in the first month maturity derivative contract set and

¹The former of these reflects the increased costs of immediate trading based on liquidity considerations, while the latter risk measures the costs associated with delayed trading at prices away from those anticipated.

²The literature based on this framework can be categorized in many ways. Perhaps the most obvious way is to label papers in terms of whether conditional or unconditional strategies are proposed. The former consists of a trade schedule that is mean-variance efficient with respect to implementation shortfall, but does not change over time to reflect changing market conditions (see, e.g., Almgren and Chriss 2001). By contrast, conditional strategies are also mean-variance efficient but adapt to dynamic variation in market conditions (see, e.g., Bertsimas and Lo 1998).

wishes to schedule trades over the period prior to the first notice day (FND) such that the position is replaced by a corresponding position in the first back month derivative contract set.³ Applying the conventional trade scheduling approach would not be appropriate. To see this, first note that the current application involves simultaneous trading in two (near) perfectly correlated price series (that is, prices of the first month and first back month derivative contract sets). Consequently, positive future shocks to prices will lead to falls in the implementation shortfall associated with the former contract set; however, the implementation shortfall associated with the latter contract set will rise and offset the former contract set implementation shortfall. The net effect means that there is essentially no timing risk in the strategy. For this reason, a different perspective on the problem is taken.

The proposed framework is one in which execution quality is optimized from the perspective of a trading desk manager who is concerned about the ex post unconditional mean and variance of execution quality over a series of trade lists (cf. the perspective of a single trader who is concerned about the ex ante moments of a single trade list).⁴ Within our framework, the mean-variance optimal trading desk manager will typically face a tradeoff between achieving a maximum level of mean execution quality with high variance (achieved by focusing trading on a particular day prior to the FND), against a lower level of mean execution quality with lower variance (achieved by spreading trading over several days).

The trader within the proposed framework could be a (stack) hedger, speculator, or arbitrageur – all that is relevant is that the trader requires a roll strategy to maintain a broader trading strategy. This includes market timing traders who seek to exploit systematic variation in roll yields (the price difference between the front month and first back month contract).⁵ This could be achieved by

³Within the context of commodity derivatives, the FND represents the day after which a roll can occur to avoid taking physical deliver of the underlying commodity.

⁴There have been previous studies of the roll decision in futures markets. Motivated by the abnormal volatility in futures prices close to their maturity (Samuelson 1965), this literature focuses on the effect of single day roll selection on the return series (and not execution quality). Using a variety of futures data, Carchano and Pardo (2009) find that there are no significant differences in the return series over a variety of different roll date selection criteria (cf. Ma et al. 1992).

⁵There is some empirical evidence that is consistent with higher roll yields prior to and during the Goldman roll period; see Mou (2011). However, many studies document contrary evidence. For instance, using different data and sample periods, Stoll and Whaley (2010) and Hamilton and Wu (2015) find little evidence of price patterns, while Bessembinder et al. (2012) document that such effects disappear within minutes. Thus it seems that as information regarding use of commercial roll strategies is in the public domain, any remaining patterns are likely to be due to limited arbitrage effects that cannot be exploited.

optimising the roll strategy with respect to roll yields as opposed to execution quality. However, the analysis presented in this paper demonstrates that such a strategy is unlikely to be highly efficient as these data do not exhibit the strong periodicities present in execution quality. Consequently, we focus on roll strategies that are optimal with respect to execution quality.⁶

The mean-variance approach adopted in this paper is analogous to the approach taken in the portfolio theory literature. Consequently, we are able to borrow concepts from this well-developed literature. In particular, the efficiency of competing roll strategies is examined using Bayesian inference based on Monte Carlo simulation; see, e.g., Kandel et al. (1995) for use of this technique within the portfolio theory literature. To anticipate some of our results, using data from five main US grain futures markets (corn, wheat, soybean, soybean meal and soybean oil), we find that roll strategies based on trading on single days prior to the FND (as is often advocated in practice) are inefficient with respect to a mean-variance optimal roll strategy. Moreover, for a sufficiently risk-averse trading desk manager, a strategy that acknowledges the need to diversify trading over several days, delivers an economically meaningful improvement in performance over a roll strategy that focuses exclusively on the maximum mean execution quality day, and a commonly used multiple-day roll strategy based on the Goldman roll dates.⁷

The remainder of this paper is organized as follows. Section 2 describes the methodologies used. Section 3 contains the application. Section 4 concludes.

2. The Framework

In this section, we introduce the measures of execution quality, formalize the roll decision faced by traders, and develop a solution to this problem. Moreover, methodologies associated with assessing the merits of this solution are presented.

2.1. Constructing the execution quality measure

The choice of which measure of execution quality to use is important. We use a transaction cost-based measure and a liquidity-based measure of execution quality, viz. the bid-ask spread and

⁶The rollover strategy is not self-financing and will require injections and withdrawals of funds because of the incurred variation in roll yields over time.

⁷See DeMiguel et al. (2009) for empirical support for the use of the $1/N$ portfolio strategy within the portfolio theory literature.

the illiquidity ratio. Specifically, the execution quality measures associated with trading the type i contract set on day n are given by

$$\text{Measure 1: } Y_{i,n} = -1 \times s_{i,n}, \quad (1a)$$

$$\text{Measure 2: } Y_{i,n} = -1 \times \sigma_{i,n}/v_{i,n}, \quad (1b)$$

where $s_{i,n}$ is the daily bid-ask spread, such that it represents the absolute value of the difference between the observed and true asset price, $\sigma_{i,n}$ is the daily volatility of the true asset price change, and $v_{i,n}$ is the daily dollar trading volume. The bid-ask spread and true price change volatility are constructed using the method of moments estimator introduced by Smith and Whaley (1994).

To concisely define this estimator, some preliminary notation is required: let day n have unit length, and let the full grid of all observation points (trades) be given by $\mathcal{G} = \{n_1, \dots, n_m\}$. Given this notation, the daily effective spread and true price change volatility are given by the solution to the following simultaneous equations:

$$\frac{1}{m} \sum_{j=1}^m |P_{i,n_{j,+}} - P_{i,n_j}| = \sqrt{\frac{2}{\pi}} \sigma_{i,n} e^{-s_{i,n}^2/2\sigma_{i,n}^2} - s_{i,n}(1 - 2\Phi(s_{i,n}/\sigma_{i,n})), \quad (2a)$$

$$\frac{1}{m} \sum_{j=1}^m |P_{i,n_{j,+}} - P_{i,n_j}|^2 = \sigma_{i,n}^2 + s_{i,n}^2 \quad (2b)$$

where P_{i,n_j} is the contract type i price associated with the j th trade of day n , $n_{j,+}$ represents the next observation after n_j on $\mathcal{G}$, and $\Phi(\cdot)$ is the standard normal cumulative density function.⁸ The estimated effective spread and true price change volatility are obtained by numerically solving these equations for each trading day and contract type in the sample.

2.2. Formalizing the Problem

The roll decision framework is designed to allow a range of unconditional roll strategies to be examined with respect to an optimal unconditional roll strategy. This framework is built on the following assumptions.

⁸See Smith and Whaley (1994) for details of the advantages of this bid-ask spread estimator over competing estimators such as Roll's (1984) serial covariance-based estimator and Thompson and Waller's (1988) mean absolute price change-based estimator.

Assumption 1. A trader is endowed with Λ contracts in the front month derivative contract set (henceforth referred to as the *type 1 set*).

Assumption 2. The trader wishes to roll the position forward to the first back month derivative contract set (henceforth the *type 2 set*). This is henceforth referred to as a *roll event*.

Assumption 3. To avoid physical deliver of the underlying asset the trader unwinds the position in the type 1 set prior to the FND and creates a new position in the type 2 set. This is henceforth referred to as the *roll strategy*.

Assumption 4. The days on which trading occurs are given by $n = \{1, 2, \dots, N\}$, with N representing the day immediately prior to the type 1 set FND (henceforth referred to as *event time*). Given this notation, the roll strategy requires that Λ contracts (henceforth the *trade target*) are held in the type 2 set on day $N + 1$.

Assumption 5. Trading can take place on one or more days prior to the FND, with $\lambda_{1,n}$ ($\lambda_{2,n}$) representing the number of type 1 (type 2) set contracts sold (bought) on the n th trading day. This implies that

$$\sum_{n=1}^N \lambda_{1,n} = \sum_{n=1}^N \lambda_{2,n} = \Lambda. \quad (3)$$

Here $\lambda_{1,n} (\geq 0)$ and $\lambda_{2,n} (\geq 0)$ are henceforth referred to as trade quantities, with the collection of trade quantities representing the *trade list*.

Remark. To minimize impact costs, this trading is assumed to take place continuously over the trading day.

Assumption 6. On each day, the type 1 set trade quantity equals the type 2 set trade quantity, thus $\lambda_n \equiv \lambda_{1,n} = \lambda_{2,n} \forall n$.

Remark. The synchronicity of the selling and buying of the type 1 and type 2 set contracts (respectively) means that the problem is reduced from $2N$ to N unknown parameters.

Assumption 7. The trading desk manager is unconditionally mean-variance efficient (MVE) with respect to the aggregate execution quality associated with the roll strategy, and requires a mean aggregate execution quality level of μ_s (standardized by portfolio size).

Remark. The unconditional mean and variance are calculated over a series of roll events.

Remark. Aggregate execution quality associated with the roll strategy is defined as follows:

$$L \equiv \sum_{n=1}^N \lambda_n Y_n, \quad \lambda_n \geq 0 \forall n, \quad (4)$$

where $Y_n = Y_{1,n} + Y_{2,n}$ is the sum of execution quality levels over the type 1 and 2 sets on day n , and $Y_{i,n}$ is the execution quality measure given by (1a) or (1b).

Under the above assumptions, the problem can be stated in terms of the following behavioral assumption. The objective of the MVE trading desk manager is to minimize the variance of aggregate execution quality (execution risk) subject to constraints (a) to (c):

$$\begin{aligned} & \underset{\boldsymbol{\lambda}}{\text{minimize}} && \boldsymbol{\lambda}^\top \boldsymbol{\Omega} \boldsymbol{\lambda} \\ & \text{subject to} && (a) \boldsymbol{\lambda}^\top \boldsymbol{\mu} = \Lambda \mu_s, \\ & && (b) \boldsymbol{1}^\top \boldsymbol{\lambda} = \Lambda, \\ & && (c) \boldsymbol{\lambda} \geq \mathbf{0}, \end{aligned}$$

where

$$\boldsymbol{\mu} = \begin{pmatrix} E[Y_1] \\ \vdots \\ E[Y_N] \end{pmatrix}, \quad \boldsymbol{\Omega} = \begin{pmatrix} \text{var}[Y_1] & \dots & \text{cov}[Y_1, Y_N] \\ \vdots & \ddots & \vdots \\ \text{cov}[Y_N, Y_1] & \dots & \text{var}[Y_N] \end{pmatrix}.$$

Here $\boldsymbol{\lambda}$ is an $(N \times 1)$ vector of trade quantities, $\boldsymbol{1}$ is an $(N \times 1)$ vector of ones, and $\mathbf{0}$ is an $(N \times 1)$ vector of zeros.

Given the definition of aggregate execution quality it follows that $\text{var}(L) = \boldsymbol{\lambda}^\top \boldsymbol{\Omega} \boldsymbol{\lambda}$, hence the objective function given in the behavioral assumption. Similarly, constraint (a) is obtained by setting $E(L) = \boldsymbol{\lambda}^\top \boldsymbol{\mu}$ to the required mean execution quality level $\Lambda \mu_s$. Constraints (b) and (c) follow from Assumption 5.

The above represents a standard quadratic programming problem, with the roll strategy based on the solution (obtained by numerical methods) referred to as the MVE roll strategy. The approach is deliberately unconditional in nature as it exploits the periodicity in execution quality observed during the period prior to the derivative contract maturity. Other financial markets do not exhibit this feature, and consequently, conditional approaches may dominate unconditional ap-

proaches in those markets. By contrast, we conjecture that the unconditional approach dominates the conditional approach in markets characterized by strong periodic variation – a view that is supported by the popularity of unconditional approaches such as the Goldman roll strategy.

2.3. Analogy to portfolio theory

Without loss of generality, portfolio size can be set equal to unity, that is, $\Lambda = 1$. In doing this, the trade quantities in λ can be interpreted as *trade weights*. Moreover, it becomes apparent that the MVE roll strategy framework is similar to the classic portfolio theory framework in which a portfolio of assets is constructed to minimize the return risk associated with a required mean return. Indeed, the above behavioral assumption is directly analogous to the case in which the variance of portfolio returns is minimized subject to a short sale constraint. Consequently, the same numerical procedures used in the portfolio theory literature can be used to derive a solution to the above optimization problem.

Within the portfolio theory framework the choice revolves around how much wealth should be invested in the N assets. By contrast, our problem consists of when trades should take place over N days. This choice delivers a minimum variance of portfolio returns or of aggregate execution quality. Moreover, this choice depends on the contemporaneous correlations among the N asset returns, or the lagged correlations among the execution quality levels observed on different days in the N periods prior to the FND. In both cases, repeated observations of these returns or execution quality levels exist over which the unconditional sample mean and variance can be estimated.

2.4. Measuring strategy performance

An examination of the quality of a particular roll strategy with respect to the above MVE roll strategy is required. To this end, we take from the portfolio theory literature and consider the standard deviation reduction achieved by conducting the MVE roll strategy with respect to the competing roll strategy. Formally, strategy inefficiency is given by

$$\rho = (\sigma_s / \sigma_*) - 1, \quad (5)$$

where $\sigma_*^2 = \lambda_*^\top \Omega \lambda_*$ is the variance of aggregate execution quality associated with the MVE strategy, $\sigma_s^2 = \lambda_s^\top \Omega \lambda_s$ is the variance of aggregate execution quality associated with the competing roll

strategy, λ_* denotes the MVE weights, λ_s denotes the weights associated with the competing roll strategy, $\lambda_*^\top \mu = \lambda_s^\top \mu = \mu_s$, $\mathbf{1}^\top \lambda_* = \mathbf{1}^\top \lambda_s = 1$, and $\lambda_* \geq \mathbf{0}$ and $\lambda_s \geq \mathbf{0}$. Note that the required mean aggregate execution quality level is dictated by the mean aggregate execution quality achieved via the competing roll strategy. Consequently, this measure gives the ratio of standard deviations for strategies with the same mean.

A graphical representation of the potential gains to using the MVE roll strategy is provided in panel (a) of Figure 1. Consider five pre-selected strategies represented by the points A, B, C, D and E. All points on the mean-variance frontier between A and E can be achieved by the MVE roll strategy, including the global minimum variance (GMV) point given by C'.

Insert Figure 1 here

The first feature of panel (a) to note is that points A and E must lie on the mean-variance frontier. These points represent the maximum and minimum mean execution quality levels, respectively. As the weights are all non-negative it follows that one cannot achieve mean execution quality levels outside of these limits. For all other points a reduction in variance is possible. It is clear that A is superior to B, B is superior to C, and so on (as we have assumed that they have equal variance). However, the variance reductions do not increase over this space. Indeed, points A and E deliver the same (zero) reduction in variance.

As an alternative we assume that the variances associated with strategies that deliver a mean execution quality level below the GMV mean execution quality level (μ_m) are given by the GMV. This amounts to replacing the minimum variance frontier from C' to E with the vertical line linking points C' to E' in panel (b). Formally, strategy inefficiency is given by

$$\rho = \begin{cases} (\sigma_s/\sigma_*) - 1, & \text{if } \mu_s \geq \mu_m, \\ (\sigma_s/\sigma_m) - 1, & \text{otherwise,} \end{cases} \quad (6)$$

where $\sigma_m^2 = \lambda_m^\top \Omega \lambda_m$ is the variance of aggregate execution quality associated with the GMV roll strategy, λ_m denotes the GMV weights, $\lambda_*^\top \mu = \lambda_s^\top \mu = \mu_s$, $\lambda_m^\top \mu = \mu_m$, $\mathbf{1}^\top \lambda_* = \mathbf{1}^\top \lambda_s = \mathbf{1}^\top \lambda_m = 1$, and $\lambda_* \geq \mathbf{0}$, $\lambda_s \geq \mathbf{0}$ and $\lambda_m \geq \mathbf{0}$. See Kandel and Stambaugh (1987) for a similar measure of inefficiency within the context of portfolio efficiency.

The second performance measure that we consider is designed to assess the economic significance of a strategy (say strategy $\mathcal{A}$) with respect to a competing strategy (say strategy $\mathcal{B}$). To this end, we consider the level of compensation δ ('performance fee' measured in terms of execution quality units) that the user of strategy $\mathcal{A}$ must receive in order to be as well off as the user of strategy $\mathcal{B}$; see Fleming et al. (2001) for use of such a performance measure within the portfolio theory literature.

Formally, a manager with the following preferences is assumed:

$$U(L) = \theta\mu_L - (1 - \theta)\sigma_L^2, \quad \theta \in [0, 1], \quad (7)$$

where $\mu_L = \boldsymbol{\lambda}^\top \boldsymbol{\mu}$, $\sigma_L^2 = \boldsymbol{\lambda}^\top \boldsymbol{\Omega} \boldsymbol{\lambda}$, and θ measures the risk preferences of the trading desk manager. Within this context, we seek δ in

$$U(L_{\mathcal{A}} + \delta) = U(L_{\mathcal{B}}). \quad (8)$$

Using (7) and this definition we can rearrange to obtain

$$\delta = \mu_{\mathcal{B}} - \mu_{\mathcal{A}} + \left(\frac{1 - \theta}{\theta} \right) (\sigma_{\mathcal{A}}^2 - \sigma_{\mathcal{B}}^2), \quad \theta \in (0, 1], \quad (9)$$

where $\mu_{\mathcal{A}} = \boldsymbol{\lambda}_{\mathcal{A}}^\top \boldsymbol{\mu}$, $\mu_{\mathcal{B}} = \boldsymbol{\lambda}_{\mathcal{B}}^\top \boldsymbol{\mu}$, $\sigma_{\mathcal{A}}^2 = \boldsymbol{\lambda}_{\mathcal{A}}^\top \boldsymbol{\Omega} \boldsymbol{\lambda}_{\mathcal{A}}$, and $\sigma_{\mathcal{B}}^2 = \boldsymbol{\lambda}_{\mathcal{B}}^\top \boldsymbol{\Omega} \boldsymbol{\lambda}_{\mathcal{B}}$.

2.5. Conducting inference

A portfolio theory literature exists that proposes methods by which one can assess whether a return portfolio is mean-variance efficient. Two primary approaches are possible: one based on classical inference and one based on Bayesian inference; see, e.g., Gibbons et al. (1989) and Kandel et al. (1995), respectively. The former tranche derives asymptotic tests (only), and is highly technical in the presence of portfolio weight restrictions; see, e.g., De Roon et al. (2001) for the case of portfolio efficiency testing when short sale constraints are imposed. To overcome these issues, we follow Li et al. (2003) and adopt an easy-to-implement Bayesian approach that is able to incorporate the finite sample uncertainty into the (posterior) distribution of the performance measures described in the previous subsection.

Our observations consist of data associated with T contract pairs (type 1 and 2 sets). Let

$t = \{1, \dots, T\}$ index the roll events over time, and $\mathbf{Y}_t$ be an $(N \times 1)$ vector containing the sum of execution quality levels over the type 1 and 2 sets for each day n in roll event t . Assume that $\mathbf{Y}_t$ has a multivariate normal distribution, with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Omega}$. We use a standard diffuse prior for this distribution

$$p(\boldsymbol{\mu}, \boldsymbol{\Omega}) \propto |\boldsymbol{\Omega}|^{-(N+1)/2}. \quad (10)$$

Given a sample of T observations, the joint posterior distribution of $\boldsymbol{\mu}$ and $\boldsymbol{\Omega}$ is given by

$$p(\boldsymbol{\mu}, \boldsymbol{\Omega} | \mathbf{Y}_1, \dots, \mathbf{Y}_T) = p(\boldsymbol{\mu} | \boldsymbol{\Omega}, \hat{\boldsymbol{\mu}}, \mathbf{Y}_1, \dots, \mathbf{Y}_T) \times p(\boldsymbol{\Omega} | \hat{\boldsymbol{\Omega}}, \mathbf{Y}_1, \dots, \mathbf{Y}_T), \quad (11)$$

where $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\Omega}}$ are the sample counterparts to $\boldsymbol{\mu}$ and $\boldsymbol{\Omega}$. Standard results demonstrate that the marginal posterior distribution $p(\boldsymbol{\Omega} | \hat{\boldsymbol{\Omega}}, \mathbf{Y}_1, \dots, \mathbf{Y}_T)$ is the inverted Wishart distribution with scale matrix $T\hat{\boldsymbol{\Omega}}$ and $T - 1$ degrees of freedom, and the conditional distribution $p(\boldsymbol{\mu} | \boldsymbol{\Omega}, \hat{\boldsymbol{\mu}}, \mathbf{Y}_1, \dots, \mathbf{Y}_T)$ is the multivariate normal distribution with mean $\hat{\boldsymbol{\mu}}$ and covariance matrix $\boldsymbol{\Omega}/T$.

The marginal posterior distributions of ρ and δ are complicated functions of the joint posterior distribution of $\boldsymbol{\mu}$ and $\boldsymbol{\Omega}$. Consequently, analytical derivations are not possible. Instead, Monte Carlo simulation is used to derive the empirical distributions based on the computed values of ρ and δ ; see Geweke (1989) for a demonstration of the accuracy of this approach. Specifically, we repeatedly draw (100000 times) a random sample of $\boldsymbol{\mu}$ and $\boldsymbol{\Omega}$ from the above posterior distributions and compute the ρ and δ values for each sample. These are then used to construct their empirical distributions upon which inference is conducted. We refer to this as the *Bayesian approach* in the remainder of the paper.

3. Application

In this section the foregoing methodologies are applied to a grain futures dataset, such that the quality of a variety of competing roll strategies are examined.

3.1. Data

All trades in the corn, wheat, soybean, soybean meal and soybean oil futures markets traded on the Chicago Mercantile Exchange (CME) over the period, January 1, 2007, to December 31,

2014, are considered.⁹ In particular, time and sales data associated with all type 1 and type 2 sets were obtained for each market (with ticker symbols CN, WC, SY, SM and BO, respectively) from *TickData, Inc.*¹⁰ These data were collected over the daytime and nighttime trading periods in which both markets are open.¹¹

The type 1 and type 2 sets are synchronized such that overlapping data over event time from 25 days before to one day before the FND (which occurs on the last business day of the month prior to the expiration month) for each roll event in the sample are considered. Missing data (or no trading days) associated with either set type result in pairwise deletion of the observation over both sets (though this occurs very rarely). A roll event coincides with the maturity of each contract in the dataset. Corn and wheat futures markets each have five maturity cycles per year, the soybean futures market has seven, and the soybean meal and soybean oil futures markets each have eight of these cycles per year. With eight years of data used, this gives rise to 40 (corn and wheat), 56 (soybean), and 64 (soybean meal and soybean oil) roll events. Moreover, as we use 25 overlapping observations per roll event this gives 1000, 1400, and 1600 daily observations, respectively.

3.2. Estimation details

For each replication of the data within the Monte Carlo simulation, the quadratic programming problem in subsection 2.2 is solved for each grain futures market using the QPROG application in GAUSS 11.0 (64-bit version). If the algorithm fails to return a converged solution then the replication is discarded and a fresh replication of the data is taken.

3.3. Data description

To provide an initial overview, the data are categorized according to the number of days to the FND, and in terms of two trading period assumptions. The first considers data from any trading period in the sample (referred to as sample A), and the second considers data from Monday to Friday daytime trading periods only (referred to as sample B). Three main characteristics of the

⁹These futures market data represent the union of the grain futures market data used in the construction of the Standard and Poor's-Goldman Sachs Commodity Index (SP-GSCI) and the Dow Jones-UBS Commodity Index (DJ-UBSCI). The latter has recently changed its title to the Bloomberg Commodity Index (BCOM).

¹⁰All prices pertain to CME Globex (electronic platform) transactions to reflect the dominance of this trading mechanism over the sample period.

¹¹The current (circa May 2015) trading times of these markets are: Sunday to Friday, 7.00pm to 7.45am (CT); and Monday to Friday, 8.30am to 1.15pm (CT).

data emerge. First, the corn futures market is the most actively traded grain futures market in the sample, while the soybean meal futures market is the least traded. Second, trading activity observed in sample B is generally higher than that observed in sample A. For this reason we make exclusive use of the latter sample in the subsequent analysis. Third, there is systematic variation (periodicity) in execution quality over event time.

The periodicity in both measures of execution quality can be appreciated by observing the shape of the mean, autocorrelation (first-order), and variance plots in panels (a), (b) and (c) of Figures 2 and 3. These are based on standardized execution quality such that for each roll event they have a zero mean and unit variance, and are plotted against event time (in days).¹² The execution quality measure for each day prior to the FND is constructed by assuming that the weight equals unity for that particular day in event time.¹³

Insert Figures 2 & 3 here

The plot in panel (a) in Figure 2 demonstrates that there is a clear pattern in mean execution quality for each grain futures market. In particular, execution quality increases up to around eight days prior to the FND and then decreases.¹⁴ This finding is consistent with the observation that contracts generally roll during this period.¹⁵ It also shows some variation in the pattern, with some markets exhibiting less defined patterns; see, e.g., corn futures execution quality periodicity. The plots in panels (b) and (c) demonstrate that there is also periodicity in autocorrelation (first-order) and variance. The same picture emerges in Figure 3.

It is also important to document the nature of the contemporaneous relations among the mea-

¹²The use of standardized execution quality is useful on two counts. First, it provides a measure of execution quality that can be used to assess its economic significance, that is, it is measured in standard deviation terms. Second, it removes any trends in execution quality over the sample period. For these reasons, standardized execution quality is used in the remainder of the paper.

¹³The mean vector and covariance matrix associated with execution quality are smoothed. In particular, the estimated mean vector and diagonal elements of the covariance matrix are obtained by using a Gaussian kernel smoother with bandwidth equal to $0.25N^{-0.4}$. For the covariance matrix we set all off-diagonal elements beyond the first-order serial covariance entries to zero.

¹⁴This inverted U-shape pattern is consistent with the results in Wang et al. (2013) who document a similar pattern in the bid-ask spreads associated with electronically traded corn futures.

¹⁵For instance, the online broker *AvaTrade* roll clients' grain futures contracts on the last Sunday prior to the FND. Moreover, though the CME does not appear to make an explicit roll date recommendation for grain futures investors, for equity index futures they impose a roll dates eight business trades prior to the FND. It should also be noted that data vendors such as Tickdata.com supply continuous futures prices under a number of different roll assumptions that include roll on the 20th calendar day of the month preceding the expiration month. This date coincides with the period six to ten days prior to the FND.

sures of execution quality. In particular, if there is a significant relationship between either of the execution quality measures and roll yields then this could invalidate the findings of the paper. For instance, a strategy that selects the best day in terms of execution quality (e.g., lowest bid-ask spread) could, at the same time, select the most undesirable day in terms of the price difference between the front month and first back month contracts (as given by the roll yield). Thus any gains in execution costs could be outweighed by losses due to the roll yield.

To investigate the above possibility the following regression model is estimated:

$$Y_{n,t} = a + bX_{n,t} + \epsilon_{n,t}, \quad n = 1, \dots, N, \quad t = 1, \dots, T, \quad (12)$$

where $Y_{n,t}$ is the execution quality measure on day n in roll event t , $X_{n,t}$ is the roll yield on day n in roll event t , and $\epsilon_{n,t}$ is a suitably defined error term. Two methodologies are used to estimate the parameters a and b . The first employs an ordinary least squares (OLS) estimator applied to pooled observations such that the parameters are assumed constant over roll event time (henceforth the pooled OLS methodology). By contrast, the second employs the methodology of Fama and MacBeth (1973) such that variation in the parameters is permitted over roll event time (henceforth the Fama-MacBeth methodology).¹⁶ Results associated with the slope coefficient, each grain market and both measures of execution quality are provided in Table 1.

Insert Table 1 here

The results provide clear evidence that a systematic relationship between execution quality and roll yields does not exist. In particular, the slope coefficient is insignificantly different from zero for almost all markets and execution quality measures. Thus the subsequent analysis pertaining to optimal decision making in terms of execution quality does not have implications for profits or losses due to price differences in the front month and first back month contracts.

¹⁶The Fama-MacBeth methodology involves a number of stages. First, the parameters are estimated within each roll event. Second, the mean and (Newey-West) standard errors of the parameters over roll event time are then constructed. Third, inference is conducted by assuming that the ratio of the mean to standard error has a normal distribution.

3.4. Assumption evaluation

The approach taken in this paper is unconditional in nature and takes advantage of the strong periodicity in execution quality evinced in Figures 2 and 3. It could be argued that a conditional approach could also be applied in which the roll decision is dependent on market conditions. One possibility is to use forecasts of execution quality within each roll event based on a dynamic model (for example, an autoregressive model) to inform this decision. The success of this approach ultimately depends on whether execution quality is predictable within this space. If the approach is successful (over all forecast horizons) then this may justify use of a conditional roll strategy. To examine this possibility the quality of forecasts based on an autoregressive model are compared to those based on a periodic model.

The framework used to examine the nature of the predictability of execution quality is built on the following forecasting model:

$$Y_{n+h,t} = a + \sum_{i=1}^p b_i Y_{n-i+1,t} + \sum_{j=1}^q c_j X_{n-j+1,t} + \sum_{k=2}^{N-h} d_k D_{k,n,t} + \epsilon_{n+h,t}, \quad n = 1, \dots, N-h, \quad t = 1, \dots, T, \quad (13)$$

where $Y_{n,t}$ is the execution quality measure on day n in roll event t , where $X_{n,t}$ is an exogenous variable observed on day n in roll event t , $D_{k,n,t}$ is a dummy variable that equals unity if k coincides with n and zero otherwise, and $\epsilon_{n+h,t}$ is a suitably defined error term. Three versions of this model are estimated. The first is a *naive model* in which $b_i = c_j = d_k = 0 \forall ijk$; the second is a *periodic model* in which $b_i = c_j = 0 \forall ij$; the third is an *autoregressive-periodic model* in which $c_k = 0 \forall k$, and the fourth is an autoregressive-periodic model that includes an exogenous variable (henceforth an *augmented autoregressive-periodic model*) in which no restrictions are imposed. The exogenous variable is given by the VIX volatility index.¹⁷

The parameters in these four models are estimated using ordinary least squares using the in-sample period, with forecasts constructed in the out-of-sample period using these parameters. Performance is assessed by constructing the mean square forecast error (MSFE) for each model with p and q set equal to one. The relative MSFE values (that is, one model's MSFE value divided

¹⁷The use of the VIX index is designed to capture the previously documented relationship between execution quality and volatility; see, e.g., Bessembinder and Seguin (1993) and Bollerslev and Melvin (1994) for empirical evidence of the relations between volatility, trading volume, market depth and bid-ask spreads.

by a competing model's MSFE value) associated with the naive, periodic, autoregressive-periodic and augmented autoregressive-periodic models are expressed in percentage terms (henceforth the *MSFE inefficiency*). These are provided in Table 2 for each grain market, various in-sample period sizes, forecast horizons of one to five days and both measures of execution quality.¹⁸

Insert Table 2 here

The results indicate that the periodic model delivers the best forecasts. For instance, when using corn futures bid-ask spread data, an in-sample size equal to half of the full sample and 1-step ahead forecasts, the MSFE inefficiency associated with the naive model with respect to the periodic model is 7.89%. By contrast, when the periodic model is compared to the autoregressive-periodic (augmented autoregressive-periodic) model, the MSFE inefficiency drops to just 1.31% (1.36%). Similar inefficiencies exist over the other futures markets, in-sample sizes and forecast horizons (particularly beyond one period). Thus, as periodic effects appear to be an important characteristic of execution quality, it follows that the unconditional approach taken in this paper is sensible.

The above forecasting models also enable examination of the implicit assumption that roll yields do not have exploitable systematic variation over event time (that is, no periodicity) – see Remark to Assumption 3. The quality of this assumption is tested by using roll yields (defined as the log of the ratio of contemporaneous type 1 to type 2 prices) as the dependent variable, and examining the quality of the naive model applied to these data with respect to the periodic model. The results in Table 2 reveal that the naive model is not inefficient in this regard. Indeed, on many occasions it is more efficient than the periodic model (that is, negative MSFE inefficiency values). This result lends support to our argument that roll yields are essentially flat over event time and can thus be excluded from the analysis of execution quality; see, e.g., Stoll and Whaley (2010) and Hamilton and Wu (2015) for consistent evidence.

3.5. *Strategy inefficiency*

To assess the performance of a particular roll strategy, information relating to the distribution of the inefficiency measure ρ obtained using the Bayesian approach described in subsection 2.4 is

¹⁸For reasons of space, results associated with the augmented autoregressive-periodic model are limited to corn futures. Use of the other futures data deliver similar results and are available on request.

provided. These inefficiencies are calculated by comparing the MVE roll strategy with a selection of s -day roll strategies in which the roll is concentrated on a single day (day s) only. The MVE roll strategy is restricted to deliver a mean aggregate execution quality level equal to that delivered by the s -day roll strategy. The results are provided in Table 3 (bid-ask spread data) and Table 4 (illiquidity ratio data).

Insert Tables 3 & 4 here

The results demonstrate that the s -day strategies are highly inefficient over all grain futures markets.¹⁹ Moreover, these inefficiencies are present over all days, though some variation is apparent. For instance, using corn futures bid-ask spread data, a roll focused on one day before the FND has a mean inefficiency of 6.70 (that is, execution quality risk is 6.70 times higher when the single-day roll strategy is used). By contrast, a roll eight days before the FND has a mean inefficiency of 1.70. It can also be observed that the inefficiencies are at their minimum value around eight days prior to the FND. Indeed, the distribution of the inefficiency measure indicates that the benefit of using the MVE roll strategy is at least zero with a posterior probability of 0.9 (or higher) around this period.²⁰

The extent of the reduction in variance available to the MVE roll strategy can be seen in panel (d) of Figures 2 and 3. Here the variances associated with this strategy are plotted against event time (in days). In contrast to the s -day roll strategy variances that are close to unity, the MVE roll strategy variance levels are considerably lower than unity, with a peak occurring around eight days prior to the FND.²¹ The variance peak coincides with the maximum mean aggregate execution quality level. Here the MVE roll strategy tends to be more concentrated (less diversified) in order to achieve this level of mean execution quality, and thus delivers a higher variance. However, some reduction in variance is possible because the location of the maximum mean aggregate execution

¹⁹As a linear transaction cost measure is considered (that is, the bid-ask spread), these results may underestimate the inefficiency associated with strategies that roll on single days. This is because such strategies may also carry a price impact cost that is not included in transaction costs.

²⁰The MVE roll strategy is characterized by a diversified portfolio, with a large number of different days used to achieve the required mean aggregate execution quality level. However, the degree of diversification decreases around the maximum mean aggregate execution quality day, with the number of different days required reaching a minimum around eight days prior to the FND. Further details regarding these characteristics are available on request.

²¹Note that the s -day and MVE strategies are equivalent in terms of mean aggregate execution quality as this is an imposed constraint.

quality day varies over the replications of the data within the simulation (which, in turn, reflects the fact that the maximum mean aggregate execution quality day varies over the sample).

3.6. Practical roll strategies

The results presented thus far provide evidence that the MVE roll strategy offers clear benefits. However, from a practical perspective, a number of challenges remain. Most importantly, it is not clear how a trader would actually implement the strategy in the absence of accurate information concerning the mean and variance of execution quality. An obvious case in point here would be the inaccuracies associated with out-of-sample estimates of these moments.

The portfolio theory literature has recently proposed alternatives to the MVE approach that avoid the need for moment estimates. In particular, DeMiguel et al. (2009) demonstrate that a naive strategy based on investing equal amounts in each asset (referred to as the $1/N$ portfolio strategy) is not outperformed by a wide range of MVE-based strategies (including one in which short sales are not permitted).

This practical method of diversification has also been adopted by commercial roll strategies (albeit with a different motivation).²² In particular, the SP-GSCI and DJ-UBSCI are “rolling indices” constructed under the assumption that commodity futures contracts roll (with uniform weights) between the fifth and ninth business days of the month preceding the expiration month (that is, the Goldman roll), and between the sixth and tenth business days of the month preceding the expiration month (that is, the UBS roll). These roll dates correspond to approximately twelve to sixteen, and eleven to fifteen days prior to the FND, respectively.²³

To examine the quality of these two approaches, we calculate the inefficiencies associated with the Goldman and UBS roll strategies with respect to the MVE roll strategy. The results in Tables 3 and 4 indicate that the UBS and Goldman roll strategies are highly inefficient with the distribution of the inefficiency measure indicating that the benefit of using the MVE roll strategy is always

²²A diversified roll is typically used in order to minimise price impact effects. However, there is an additional benefit in terms of execution quality risk reduction – a feature that is the focus of the current paper.

²³The Goldman and UBS rolls are commonly used in practice. For instance, there are a number of exchange traded funds (ETFs) that seek to track the SP-GSCI and DJ-UBSCI – each adopting the Goldman or UBS roll, respectively. To give an idea of the scale of this activity it should be noted that the *iShares S&P GSCI Commodity-Indexed Trust* has a market capitalization of \$760,440,447 on September 10, 2015. This information was taken from the iShares website on <https://www.ishares.com/us/products/239757/ishares-sp-gsci-commodityindexed-trust-fund> accessed on September 11, 2015.

above zero even with a posterior probability of 0.99.

The primary reason for the failure of the Goldman and UBS roll strategies is the location of days on which the roll occurs with respect to the maximum execution quality day (see Figures 2 and 3). To address this shortcoming, we propose a multiple-day roll strategy in which the trader equally spreads the roll centered on the maximum execution quality day. Two variants are considered: strategies with weights equal to unity on the center point (the *adaptive single-day strategy*), and weights equal to one fifth on each day around (and including) the center point (the *adaptive multiple-day strategy*). In both cases, to ensure real time consistency the maximum execution quality day is estimated using an in-sample period, with the quality of the strategies assessed using an out-of-sample period (ergo the ‘adaptive’ adjective). In the subsequent analysis we use the first ten sets of event time as the in-sample period, and the remaining sets of event time as the out-of-sample period.

Use of these strategies will ultimately lead to a mean-variance tradeoff. When a roll occurs exclusively on the center point, a high mean and variance is obtained. As the degree of roll diversification increases, a lower mean *and* variance is obtained. Thus, this provides a choice to the trading desk manager that they can optimize based on their risk preferences. For a manager that seeks a high mean execution quality level (low risk-averse manager), a concentrated roll is suitable. By contrast, for a manager seeking a low variance execution quality level (high risk-averse manager), a diversified roll is best.

There is an addition benefit from using the diversified strategy in that the estimation error associated with the location of the roll is reduced. By spreading the roll one is more likely to capture the true maximum execution quality day. This will ultimately lead to a lower mean performance associated with the single-day strategy, such that it may be dominated by multiple-day strategy in terms of mean and variance.

3.7. Economic significance

To assess the economic significance of two competing strategies (strategy $\mathcal{A}$ with respect to strategy $\mathcal{B}$), the distribution of δ obtained using the Bayesian approach is provided in Table 5. Here strategy $\mathcal{A}$ represents the adaptive single-day roll strategy, and strategy $\mathcal{B}$ represents the adaptive multiple-day roll strategy. We consider $\theta = \{0.1, 0.3, 0.5, 0.7, 0.9\}$. As execution quality

is standardized, it follows that the performance fee is given in terms of execution quality standard deviations. Space limitations motivate presentation of results based on bid-ask spread data only.²⁴

Insert Table 5 here

The results indicate that for trading desk managers with a mean focus (that is, $\theta = 0.9$), the mean performance fee is close to zero. That is, trading desk managers using strategy $\mathcal{A}$ do not require a significant amount to be as well off as strategy $\mathcal{B}$ users. As the preference for mean switches toward a preference for variance (that is, θ falls), the performance fee increases and becomes economically significant in terms of magnitude. Importantly, the performance fees are at least positive with a posterior probability of 0.9, when θ falls below 0.5. Thus, for a sufficiently risk-averse trading desk manager, the adaptive multiple-day roll strategy has (economic) value.

It is possible to use this framework to examine the relative quality of competing multiple-day strategies. In particular, we calculate the performance fees associated with the Goldman roll strategy (strategy $\mathcal{A}$) and the adaptive multiple-day roll strategy (strategy $\mathcal{B}$).²⁵ Results are provided in Table 6.

Insert Table 6 here

The results demonstrate that the mean performance fees are generally positive and economically meaningful in magnitude. Thus, traders using the Goldman roll strategy require an economically significant level of compensation. Moreover, for mean-focused trading desk managers these performance fees are at least positive with a posterior probability of 0.9. This result is consistent with the observation that the Goldman roll strategy is centered on days with execution quality levels that are significantly below the maximum available.²⁶

²⁴The assumption is maintained in the subsequent analysis. Results pertaining to the use of illiquidity ratio data give similar results. These are available on request.

²⁵Similar results are obtained when the Goldman roll strategy is replaced by the UBS roll strategy. These are available on request.

²⁶The Goldman roll strategy appears to improve for variance-focused trading desk managers (all markets except wheat). As both strategies are based on the same number of roll days (that is, five), and the variances of execution quality are broadly similar over the extant roll days, then this result is due to lower first-order autocorrelations observed during the time in which the Goldman roll occurs (see panel (b) of Figure 2).

3.8. Intraday variation

The analysis thus far has assumed that traders roll their positions over the entire trading day. To examine this supposition we repeat the analysis under the assumption that trading occurs during more specific periods within the trading day. In particular, three intraday periods are considered: viz., the early trading session (8.30am to 10.00am), the mid trading session (10.01am to 12.00pm), and the late trading session (12.01pm to 1.15pm). The mean, autocorrelation (first-order) and variances associated with the single-day and MVE strategies applied to corn futures data are plotted in panels (a), (b), (c) and (d) of Figure 4. In addition, the distribution of the performance fee associated with corn futures, the adaptive single-day roll strategy (strategy $\mathcal{A}$), and the adaptive multiple-day roll strategy (strategy $\mathcal{B}$), are summarized in Table 7.²⁷

Insert Figure 4 here

Insert Table 7 here

The plots in Figure 4 indicate that the means, autocorrelations and variances are similar across the three intraday trading sessions, and are comparable to the plots associated with the full day trading session. Moreover, the results in Table 7 indicate that performance fees are above zero with posterior probability 0.95 for θ levels below 0.5 – a result that is similar to that observed previously. Thus, the results appear robust to variation in the time at which the trader rolls their position within the trading day.

3.9. The Goldman roll revisited

The documented poor performance associated with the Goldman roll strategy, particularly compared to the adaptive multiple-day roll strategy (see Table 6), requires further analysis. As the degree of diversification across these strategies is equal, then the performance difference must be due to the location of the roll in event time. Specifically, the Goldman roll occurs too early. It could be argued that this result is specific to grain futures markets with FNDs that occur late in the month preceding the expiration month. However, there is nothing unusual about this timing as many other commodity futures markets share similar FND dates. Indeed, for energy futures

²⁷Results pertaining to the other grain futures markets give similar results. These are available on request.

such as heating oil futures (HO), the FND is the second business day *after* the last business day of the month preceding the expiration month. Thus, for many commonly used commodity futures markets the Goldman roll is likely to be mistimed.

The question that naturally arises from the above discussion is: are there any futures market that are likely to be suited to the Goldman roll? The Goldman roll is centered on the seventh business day of the month preceding the expiration month. Given the evidence presented above that execution quality peaks around eight business days prior to the FND (see Figure 2), it follows that we seek a futures market with an FND (or last trading day, LTD, if this is prior to the FND) that occurs around the fifteenth business day of the month preceding the expiration month. With approximately twenty one business days in a month, then would like to consider a futures market with an FND (or LTD) that occurs six business days before the last business day of the month preceding the expiration month. Coffee and crude oil futures (KC and CL) traded on the Intercontinental Exchange (ICE) and CME, respectively, possess this feature and are likely to be suited to the Goldman roll.

To investigate the above conjecture we consider the Goldman roll within the context of the coffee and crude oil futures markets. This examination is based on using coffee and crude oil futures bid-ask spread data over the same period as used in the preceding analysis.²⁸ The results in Figure 5 provide the mean, autocorrelation (first-order) and variances plots associated with coffee futures execution quality over event time. It is apparent that the peak in mean execution quality now occurs around the same time as the Goldman roll (cf. the plots in Figures 2, 3 and 4). This is simply because the FND is closer to the Goldman roll dates.

Insert Figure 5 here

To further examine the above issue, performance fees associated with the Goldman roll with respect to the adaptive single-day and multiple-day roll strategies are calculated using coffee and crude oil futures bid-ask spread data. The results in Table 8 indicate that the Goldman roll is no worse than the competing strategies. This is in contrast to the results associated with other futures markets in which the Goldman roll is notably inferior to the adaptive multiple-day strategy for

²⁸Bid-ask spreads are calculated using time and sales coffee and crude oil futures data obtained from *TickData, Inc.*.

mean-focused traders (see Table 6). This evidence supports the hypothesis that the Goldman roll appears to work efficiently for only a handful of (albeit important) futures markets.

Insert Table 8 here

4. Conclusions

In this paper a mean-variance framework applicable to traders' who roll their positions in securities from short to longer-dated contracts is proposed. The empirical findings are summarized as follows:

1. There is a distinct inverted U-shaped pattern in execution quality in the 25-day period prior to the FND, with a peak mean execution quality level observed around eight days prior to the FND.
2. Rolls based on single days prior to the FND are likely to be inefficient with respect to the MVE strategy. One possible exception would be those strategies that roll close to the maximum observed mean execution quality levels.
3. The performance of the one-size-fits-all Goldman and UBS roll strategies is generally poor in comparison to competing roll strategies (though exemptions are documented) . This is primarily because the former strategies are centered on days with mean execution quality levels that are below the maximum available.
4. An adaptive multiple-day roll strategy that evenly spreads the roll around the maximum mean execution quality day delivers lower mean execution quality levels, but with less execution risk, in comparison to an adaptive single-day roll strategy. Indeed, for sufficiently risk-averse traders, the adaptive multiple-day roll strategy is shown to be preferable.

These results have obvious implications for a variety of participants in financial markets. First, for a trading desk manager overseeing traders with roll requirements (in turn under instruction from a portfolio manager or stack hedger), it is not advisable to undertake a concentrated roll on single days. Rather, a diversified roll strategy should be adopted (providing they are reasonably risk-averse with respect to execution risk). Second, the results have implications for researchers and financial institutions who work in the area of index construction. For instance, the SP-GSCI

is constructed by implementing the Goldman roll. Consequently, those seeking to track this index are likely to be exposed to too high a level of execution risk.

References

- Almgren, R., Chriss, N., 2001. Optimal execution of portfolio transactions. *Journal of Risk* 3, 5–39.
- Bertsimas, D., Lo, A., 1998. Optimal control of liquidation costs. *Journal of Financial Markets* 1, 1–50.
- Bessembinder, H., Carrion, A., Tuttle, L., Venkataraman, K., 2012. Predatory or sunshine trading? Evidence from crude oil ETF rolls. Southern Methodist University Working Paper.
- Bessembinder, H., Seguin, P., 1993. Price volatility, trading volume, and market depth: Evidence from futures markets. *Journal of Financial and Quantitative Analysis* 28, 21–39.
- Bollerslev, T., Melvin, M., 1994. Bid-ask spreads and volatility in the foreign exchange market: An empirical analysis. *Journal of International Economics* 36, 355–372.
- Carchano, Ó., Pardo, Á., 2009. Rolling over stock index futures contracts. *Journal of Futures Markets* 29, 684–694.
- DeMiguel, V., Garlappi, L., Uppal, R., 2009. Optimal versus naive diversification: How inefficient is the $1/N$ portfolio strategy? *Review of Financial Studies* 22, 1915–1953.
- De Roon, F., Nijman, T., Werker B., 2001. Testing for mean-variance spanning with short sales constraints and transaction costs: The case of emerging markets. *Journal of Finance* 61, 721–742.
- Engle, R., Ferstenberg, R., 2007. Execution risk. *Journal of Trading* 2, 10–20.
- Fama, E., MacBeth, J., 1973. Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy* 81, 607–636.
- Fleming, J., Kirby, C., Ostdiek, B., 2001. The economic value of volatility timing. *Journal of Finance* 56, 329–352.
- Forsyth, P., Kennedy, J., Tse, S., Windcliff, H., 2012. Optimal trade execution: A mean quadratic variation approach. *Journal of Economic Dynamics and Control* 36, 1971–1991.

- Geweke, J., 1989. Bayesian inference in econometric models using Monte Carlo integration. *Econometrica* 57, 1317–1339.
- Gibbons, M., Ross, S., Shanken, L., 1989. A test of the efficiency of a given portfolio. *Econometrica* 57, 1121–1152.
- Hamilton, J., Wu, J., 2015. Effects of index-fund investing on commodity futures prices. *International Economic Review* 56, 187–205.
- He, H., Mamaysky, H., 2005. Dynamic trading policies with price impact. *Journal of Economic Dynamics and Control* 29, 891–930.
- Kandel, S., Stambaugh, R., 1987. On correlations and inferences about mean-variance efficiency. *Journal of Financial Economics* 18, 61–90.
- Kandel, S., McCulloch, R., Stambaugh, R., 1995. Bayesian inference and portfolio efficiency. *Review of Financial Studies* 8, 1–53.
- Li, K., Sarkar, A., Wang, Z., 2003. Diversification benefits of emerging markets subject to portfolio constraints. *Journal of Empirical Finance* 10, 57–80.
- Ma, K., Mercer, M., Walker, M., 1992. Rolling over futures contracts: A note. *Journal of Futures Markets* 12, 203–217.
- Mou, Y., 2011. Limits to arbitrage and commodity index investment: Front-running the Goldman roll. SSRN Working Paper available at <http://ssrn.com/abstract=1716841>.
- Perold, A., 1988. The implementation shortfall: Paper versus reality. *Journal of Portfolio Management* 14, 4–9.
- Roll, R., 1984. A simple implicit measure of bid/ask spread in an efficient market. *Journal of Finance* 39, 1127–1139.
- Samuelson, P., 1965. Proof that properly anticipated prices fluctuate randomly. *Industrial Management Review* 6, 41–49.
- Schied, A., Schoeneborn, T., 2009. Risk aversion and the dynamics of optimal liquidation strategies in illiquid markets. *Finance and Stochastics* 13, 181–204.

- Smith, T., Whaley, R., 1994. Estimating the effective bid/ask spread from time and sales data. *Journal of Futures Markets* 14, 437–455.
- Stoll, H., Whaley, R., 2010. Commodity index investing and commodity futures prices. *Journal of Applied Finance* 20, 7–46.
- Thompson, S., Waller, M., 1988. Determinants of liquidity costs in commodity futures markets. *Review of Futures Markets* 7, 110–126.
- Tse, S., Forsyth, P., Kennedy, J., Windcliff, H., 2013. Comparison between the mean-variance optimal and the mean-quadratic-variation optimal trading strategies. *Applied Mathematical Finance* 20, 415–449.
- Wang, X., Garcia, P., Irwin, S., 2013. The behavior of bid-ask spreads in the electronically-traded corn futures market. *American Journal of Agricultural Economics* 96, 557–577.

Table 1 – The contemporaneous relations between execution quality and roll yield

Methodology	Dependent Variable	Statistic			
		coeff.	s.e.	z-stat.	p-value
(a) Corn Data					
Pooled OLS	Bid-Ask Spread	0.02	0.04	0.59	0.28
	Illiquidity Ratio	0.03	0.04	0.99	0.16
Fama-MacBeth	Bid-Ask Spread	0.02	0.06	0.37	0.36
	Illiquidity Ratio	0.03	0.07	0.48	0.31
(b) Wheat Data					
Pooled OLS	Bid-Ask Spread	−0.02	0.04	−0.68	0.25
	Illiquidity Ratio	0.02	0.03	0.70	0.24
Fama-MacBeth	Bid-Ask Spread	−0.02	0.07	−0.36	0.36
	Illiquidity Ratio	0.02	0.06	0.39	0.35
(c) Soybean Data					
Pooled OLS	Bid-Ask Spread	0.01	0.03	0.24	0.41
	Illiquidity Ratio	0.03	0.03	1.01	0.16
Fama-MacBeth	Bid-Ask Spread	0.01	0.05	0.15	0.44
	Illiquidity Ratio	0.03	0.06	0.56	0.29
(d) Soybean Meal Data					
Pooled OLS	Bid-Ask Spread	−0.07	0.03	−2.56	0.01
	Illiquidity Ratio	−0.02	0.03	−0.87	0.19
Fama-MacBeth	Bid-Ask Spread	−0.07	0.05	−1.55	0.06
	Illiquidity Ratio	−0.02	0.05	−0.46	0.32
(e) Soybean Oil Data					
Pooled OLS	Bid-Ask Spread	0.03	0.03	1.04	0.15
	Illiquidity Ratio	0.01	0.03	0.18	0.43
Fama-MacBeth	Bid-Ask Spread	0.03	0.04	0.93	0.18
	Illiquidity Ratio	0.01	0.04	0.15	0.44

Notes: This table contains the estimated (mean) slope coefficient, standard error, z-statistic and p-value associated with a regression of each measure of execution quality on roll yield. The pooled OLS and Fama-MacBeth methodologies are used to estimate these statistics with a Newey-West estimator employed to calculate the standard error in the case of the latter methodology. The p-value examines the null hypothesis that the coefficient equals zero.

Table 2 – MSFE inefficiencies (autoregressive and periodic effects)

Variable	Model Comparison	Forecast Horizon	In-sample Size (w.r.t. full sample)				
			30%	40%	50%	60%	70%
(a) Corn Data							
Bid-Ask Spread	M0 v. M1	All	3.15	3.30	7.89	11.76	11.44
		M1 v. M2	1	−1.03	1.20	1.31	1.07
	2		0.20	0.33	0.28	0.12	−0.07
	3		−0.16	−0.10	−0.08	0.01	0.02
	4		−0.87	−0.20	0.06	−0.17	−0.42
	5		−0.45	−0.21	−0.07	0.13	0.15
	M1 v. M3	1	−2.82	0.12	1.36	1.10	1.21
		2	−1.05	0.34	0.43	0.28	−0.02
		3	−3.75	−0.36	0.06	0.14	−0.10
		4	−5.68	−0.52	0.07	−0.17	−0.44
5		−5.01	−0.52	−0.07	0.16	0.11	
Illiquidity Ratio	M0 v. M1	All	7.98	10.24	11.10	10.66	7.92
Roll Yield	M0 v. M1	All	−3.08	−0.20	0.19	0.24	2.12
(b) Wheat Data							
Bid-Ask Spread	M0 v. M1	All	19.75	25.92	31.46	30.00	27.86
	M1 v. M2	1	8.24	8.03	4.21	6.92	7.35
		2	−0.44	−0.12	0.27	0.36	−0.20
		3	−2.91	−1.31	−1.38	1.20	0.60
		4	−0.58	0.41	1.25	1.88	2.38
		5	−0.74	−1.59	−0.10	0.15	−0.45
Illiquidity Ratio	M0 v. M1	All	51.91	61.83	72.09	73.66	77.37
Roll Yield	M0 v. M1	All	−14.63	−3.81	−0.17	4.78	2.50
(c) Soybean Data							
Bid-Ask Spread	M0 v. M1	All	6.21	5.50	7.29	8.82	4.58
	M1 v. M2	1	6.61	6.91	8.97	11.26	10.52
		2	4.14	3.12	4.82	4.36	5.68
		3	1.17	1.37	1.35	1.23	1.51
		4	0.31	0.20	0.06	0.21	0.87
		5	0.09	0.12	−0.22	−0.12	0.12
Illiquidity Ratio	M0 v. M1	All	19.01	19.89	23.75	24.92	20.98
Roll Yield	M0 v. M1	All	−5.00	−6.24	−1.45	1.56	0.42
(d) Soybean Meal Data							
Bid-Ask Spread	M0 v. M1	All	6.80	3.01	7.02	8.00	8.19
	M1 v. M2	1	5.64	7.37	5.59	6.68	2.27
		2	1.86	2.68	2.02	1.42	0.21
		3	0.90	1.00	0.64	0.07	−1.11
		4	−0.05	−0.21	−0.15	0.00	−0.14
		5	−0.30	0.03	0.00	−0.07	0.06
Illiquidity Ratio	M0 v. M1	All	12.44	10.76	13.26	16.09	13.98
Roll Yield	M0 v. M1	All	−3.53	4.43	5.15	−1.59	−2.63
(e) Soybean Oil Data							
Bid-Ask Spread	M0 v. M1	All	9.37	13.41	17.00	16.69	12.05
	M1 v. M2	1	4.04	6.29	4.49	3.26	2.68
		2	0.77	1.21	0.56	0.26	0.74
		3	0.65	0.91	0.54	0.18	−0.33
		4	0.05	0.03	−0.17	−0.41	−0.25
		5	0.13	0.04	−0.25	−0.08	−0.06
Illiquidity Ratio	M0 v. M1	All	19.78	20.82	24.50	29.30	30.56
Roll Yield	M0 v. M1	All	−12.90	−9.26	−2.97	−5.39	−2.44

Notes: This table contains the out-of-sample MSFE inefficiencies associated with the naive (M0), periodic (M1), autoregressive-periodic (M2) and augmented autoregressive-periodic (M3) models. These are provided for bid-ask spread, illiquidity ratio and roll yield forecasts based on various in-sample period sizes and forecast horizons. Inefficiencies are given in percentage terms.

Table 3 – Strategy inefficiency (bid-ask spread data)

Strategy	Statistic							
	MN	SD	IQR	50%	75%	90%	95%	99%
(a) Corn Data								
1 (day(s) to FND)	6.70	1.86	2.31	6.45	5.40	4.59	4.15	3.41
4	4.93	1.94	2.45	4.83	3.64	2.55	1.93	0.84
8	1.70	1.70	1.90	1.25	0.49	0.00	0.00	0.00
12	6.37	2.01	2.45	6.19	5.06	4.10	3.46	1.97
16	5.27	1.69	2.07	5.11	4.15	3.34	2.84	1.72
24	7.11	2.00	2.49	6.85	5.73	4.85	4.35	3.39
Goldman Roll	2.06	0.71	0.89	1.96	1.55	1.26	1.10	0.84
UBS Roll	2.15	0.72	0.91	2.05	1.64	1.32	1.16	0.89
(b) Wheat Data								
1	8.34	2.28	2.86	8.02	6.73	5.76	5.24	4.36
4	6.10	1.88	2.33	5.91	4.84	3.97	3.46	2.33
8	1.06	1.04	1.21	0.83	0.30	0.00	0.00	0.00
12	4.48	1.64	2.06	4.32	3.36	2.58	2.13	1.28
16	4.99	1.43	1.78	4.79	3.99	3.36	3.03	2.47
24	6.45	1.80	2.26	6.21	5.18	4.41	4.00	3.31
Goldman Roll	2.10	0.69	0.86	2.00	1.62	1.33	1.18	0.92
UBS Roll	1.99	0.67	0.83	1.89	1.52	1.24	1.09	0.84
(c) Soybean Data								
1	7.52	1.44	1.88	7.38	6.51	5.80	5.41	4.74
4	4.19	1.40	1.86	4.15	3.23	2.46	2.01	1.08
8	0.94	0.64	0.79	0.86	0.50	0.00	0.00	0.00
12	3.85	0.83	1.09	3.78	3.27	2.86	2.62	2.19
16	4.19	0.86	1.12	4.11	3.59	3.16	2.93	2.53
24	5.22	1.04	1.35	5.12	4.49	3.98	3.71	3.22
Goldman Roll	1.43	0.36	0.47	1.39	1.17	1.00	0.90	0.74
UBS Roll	1.40	0.36	0.47	1.37	1.15	0.98	0.88	0.73
(d) Soybean Meal Data								
1	5.46	0.97	1.28	5.38	4.78	4.29	4.01	3.53
4	3.53	1.23	1.65	3.50	2.69	1.98	1.58	0.72
8	1.12	0.79	1.00	1.03	0.57	0.00	0.00	0.00
12	3.74	0.88	1.14	3.71	3.15	2.65	2.35	1.76
16	3.90	0.74	0.97	3.85	3.39	3.01	2.80	2.41
24	4.63	0.84	1.11	4.56	4.04	3.61	3.38	2.96
Goldman Roll	1.37	0.32	0.41	1.34	1.14	0.99	0.90	0.75
UBS Roll	1.33	0.32	0.41	1.30	1.11	0.95	0.87	0.72
(e) Soybean Oil Data								
1	5.96	1.05	1.37	5.87	5.22	4.69	4.39	3.87
4	3.99	1.05	1.36	3.99	3.31	2.69	2.31	1.55
8	0.97	0.72	0.95	0.88	0.44	0.00	0.00	0.00
12	3.53	0.92	1.20	3.51	2.92	2.40	2.09	1.50
16	3.77	0.72	0.94	3.72	3.27	2.91	2.70	2.33
24	5.46	0.98	1.28	5.38	4.77	4.29	4.01	3.53
Goldman Roll	1.37	0.32	0.42	1.34	1.14	0.98	0.90	0.75
UBS Roll	1.32	0.32	0.42	1.29	1.09	0.93	0.85	0.70

Notes: This table contains information relating to the distribution of the inefficiency measure ρ obtained using the Bayesian approach applied to bid-ask spread data. Specifically, the mean (MN), standard deviation (SD), interquartile range (IR), and 50, 75, 90, 95, and 99 percentiles are provided. These inefficiencies are based on a comparison of the MVE roll strategy with the s -day roll strategy in which the roll takes place on one day (day s) only, and the multiple-day Goldman and UBS roll strategies.

Table 4 – Strategy inefficiency (illiquidity ratio data)

Strategy	Statistic							
	MN	SD	IQR	50%	75%	90%	95%	99%
(a) Corn Data								
1 (day(s) to FND)	9.15	2.46	3.07	8.81	7.42	6.36	5.80	4.86
4	6.67	2.41	2.94	6.53	5.13	3.81	2.93	1.11
8	1.70	1.57	1.91	1.34	0.54	0.00	0.00	0.00
12	3.11	1.26	1.57	2.95	2.24	1.68	1.36	0.77
16	5.25	1.49	1.86	5.05	4.21	3.56	3.22	2.66
24	8.44	2.27	2.85	8.13	6.84	5.87	5.34	4.48
Goldman Roll	1.76	0.61	0.76	1.67	1.33	1.08	0.95	0.74
UBS Roll	1.42	0.54	0.66	1.34	1.04	0.83	0.71	0.53
(b) Wheat Data								
1	13.88	3.66	4.55	13.37	11.31	9.74	8.94	7.55
4	8.88	2.41	3.01	8.55	7.18	6.15	5.60	4.69
8	1.49	1.23	1.57	1.30	0.58	0.00	0.00	0.00
12	2.89	1.27	1.56	2.73	2.03	1.47	1.13	0.07
16	5.23	1.48	1.86	5.02	4.19	3.55	3.21	2.65
24	11.39	3.03	3.78	10.95	9.25	7.96	7.29	6.15
Goldman Roll	2.00	0.68	0.84	1.90	1.52	1.24	1.09	0.85
UBS Roll	1.62	0.60	0.73	1.53	1.20	0.95	0.82	0.61
(c) Soybean Data								
1	10.24	1.91	2.48	10.06	8.90	7.97	7.46	6.59
4	6.34	1.49	1.87	6.28	5.39	4.59	4.06	2.81
8	0.84	0.62	0.80	0.77	0.40	0.00	0.00	0.00
12	3.05	0.74	0.96	2.99	2.54	2.18	1.97	1.61
16	4.53	0.93	1.21	4.45	3.88	3.43	3.18	2.76
24	8.06	1.53	2.00	7.91	6.98	6.24	5.83	5.14
Goldman Roll	1.48	0.37	0.49	1.44	1.21	1.04	0.94	0.77
UBS Roll	1.27	0.34	0.44	1.24	1.04	0.87	0.79	0.64
(d) Soybean Meal Data								
1	6.70	1.15	1.51	6.60	5.89	5.31	4.99	4.42
4	4.40	1.19	1.50	4.38	3.65	2.97	2.54	1.52
8	1.08	0.71	0.95	1.03	0.58	0.00	0.00	0.00
12	3.33	0.72	0.94	3.28	2.83	2.46	2.24	1.87
16	4.69	0.84	1.11	4.62	4.09	3.67	3.44	3.03
24	7.28	1.24	1.63	7.18	6.41	5.78	5.44	4.84
Goldman Roll	1.51	0.34	0.44	1.48	1.27	1.10	1.01	0.85
UBS Roll	1.33	0.31	0.41	1.31	1.11	0.96	0.87	0.73
(e) Soybean Oil Data								
1	8.56	1.43	1.88	8.44	7.56	6.82	6.43	5.74
4	5.78	1.07	1.38	5.71	5.05	4.50	4.19	3.52
8	1.22	0.88	1.21	1.15	0.57	0.00	0.00	0.00
12	2.51	0.74	0.96	2.47	2.01	1.61	1.37	0.85
16	4.04	0.74	0.97	3.98	3.52	3.15	2.94	2.59
24	8.58	1.44	1.89	8.46	7.57	6.85	6.45	5.75
Goldman Roll	1.21	0.29	0.39	1.19	1.00	0.86	0.78	0.65
UBS Roll	1.01	0.27	0.35	0.98	0.82	0.69	0.62	0.50

Notes: This table contains information relating to the distribution of the inefficiency measure ρ obtained using the Bayesian approach applied to illiquidity ratio data. Specifically, the mean (MN), standard deviation (SD), interquartile range (IR), and 50, 75, 90, 95, and 99 percentiles are provided. These inefficiencies are based on a comparison of the MVE roll strategy with the s -day roll strategy in which the roll takes place on one day (day s) only, and the multiple-day Goldman and UBS roll strategies.

Table 5 – Performance fees (adaptive single-day roll strategy v. adaptive multiple-day roll strategy)

Preference ($1 - \theta$)	Statistic							
	MN	SD	IQR	5%	10%	50%	90%	95%
(a) Corn Data								
0.1 (mean-focused)	0.26	0.24	0.20	0.58	0.47	0.24	0.05	-0.01
0.3	0.44	0.61	0.29	1.08	0.79	0.34	0.12	0.05
0.5	0.78	1.76	0.54	2.19	1.49	0.52	0.18	0.07
0.7	1.54	4.43	1.17	4.75	3.13	0.94	0.25	0.00
0.9 (variance-focused)	5.47	17.36	4.45	17.62	11.48	3.08	0.48	-0.44
(b) Wheat Data								
0.1	-0.01	0.14	0.10	0.16	0.10	-0.02	-0.11	-0.15
0.3	0.08	0.69	0.14	0.40	0.25	0.03	-0.07	-0.11
0.5	0.25	0.76	0.26	0.89	0.58	0.12	-0.04	-0.09
0.7	0.63	2.29	0.56	2.15	1.37	0.33	0.00	-0.10
0.9	2.64	23.93	2.13	8.31	5.43	1.37	0.17	-0.20
(c) Soybean Data								
0.1	0.13	0.14	0.18	0.37	0.31	0.13	-0.04	-0.09
0.3	0.34	0.18	0.22	0.66	0.57	0.32	0.13	0.08
0.5	0.71	0.31	0.37	1.27	1.10	0.66	0.38	0.31
0.7	1.58	0.65	0.77	2.79	2.41	1.46	0.89	0.76
0.9	5.92	2.50	2.89	10.53	9.04	5.46	3.32	2.87
(d) Soybean Meal Data								
0.1	0.15	0.14	0.19	0.39	0.33	0.15	-0.03	-0.08
0.3	0.37	0.18	0.23	0.69	0.60	0.36	0.16	0.11
0.5	0.77	0.30	0.37	1.31	1.15	0.74	0.43	0.36
0.7	1.71	0.63	0.77	2.86	2.52	1.62	1.00	0.86
0.9	6.37	2.39	2.90	10.75	9.42	6.00	3.74	3.23
(e) Soybean Oil Data								
0.1	0.06	0.10	0.13	0.22	0.18	0.06	-0.06	-0.10
0.3	0.20	0.12	0.15	0.41	0.36	0.19	0.05	0.02
0.5	0.46	0.20	0.25	0.81	0.71	0.43	0.23	0.18
0.7	1.06	0.42	0.51	1.83	1.60	1.00	0.59	0.50
0.9	4.06	1.57	1.91	6.94	6.09	3.81	2.32	1.97

Notes: This table contains information relating to the distribution of the performance fee δ obtained using the Bayesian approach. Specifically, the mean (MN), standard deviation (SD), interquartile range (IR), and 5, 10, 50, 90, and 95 percentiles are provided. The performance fees are based on a comparison of the adaptive single-day roll strategy and the adaptive multiple-day roll strategy.

Table 6 – Performance fees (Goldman roll strategy v. adaptive multiple-day roll strategy)

Preference ($1 - \theta$)	Statistic							
	MN	SD	IQR	5%	10%	50%	90%	95%
(a) Corn Data								
0.1 (mean-focused)	0.20	0.12	0.12	0.38	0.33	0.21	0.08	0.03
0.3	0.19	0.22	0.14	0.40	0.34	0.20	0.04	−0.04
0.5	0.16	0.44	0.19	0.50	0.39	0.18	−0.07	−0.24
0.7	0.10	1.56	0.35	0.81	0.54	0.15	−0.38	−0.75
0.9 (variance-focused)	−0.24	4.41	1.27	2.41	1.42	−0.02	−2.02	−3.43
(b) Wheat Data								
0.1	0.50	0.11	0.11	0.65	0.61	0.50	0.38	0.34
0.3	0.50	0.21	0.13	0.70	0.64	0.50	0.37	0.31
0.5	0.51	0.37	0.17	0.84	0.72	0.50	0.31	0.20
0.7	0.54	0.79	0.30	1.23	0.94	0.52	0.15	−0.10
0.9	0.65	4.13	1.07	3.30	2.16	0.57	−0.79	−1.75
(c) Soybean Data								
0.1	0.33	0.09	0.12	0.48	0.45	0.33	0.21	0.17
0.3	0.28	0.10	0.13	0.44	0.41	0.29	0.15	0.11
0.5	0.20	0.14	0.17	0.40	0.36	0.21	0.02	−0.05
0.7	−0.01	0.26	0.31	0.35	0.28	0.02	−0.33	−0.47
0.9	−1.01	0.94	1.09	0.26	0.00	−0.89	−2.18	−2.69
(d) Soybean Meal Data								
0.1	0.25	0.09	0.12	0.40	0.37	0.25	0.14	0.11
0.3	0.22	0.10	0.13	0.38	0.34	0.22	0.10	0.06
0.5	0.16	0.13	0.16	0.35	0.31	0.16	0.00	−0.06
0.7	0.01	0.23	0.28	0.35	0.27	0.02	−0.27	−0.38
0.9	−0.73	0.80	0.97	0.44	0.19	−0.66	−1.73	−2.13
(e) Soybean Oil Data								
0.1	0.29	0.08	0.11	0.42	0.39	0.29	0.18	0.15
0.3	0.27	0.09	0.11	0.41	0.38	0.27	0.16	0.13
0.5	0.24	0.11	0.14	0.41	0.37	0.24	0.10	0.06
0.7	0.17	0.19	0.23	0.46	0.39	0.18	−0.06	−0.15
0.9	−0.18	0.66	0.79	0.82	0.59	−0.15	−0.99	−1.30

Notes: This table contains information relating to the distribution of the performance fee δ obtained using the Bayesian approach. Specifically, the mean (MN), standard deviation (SD), interquartile range (IR), and 5, 10, 50, 90, and 95 percentiles are provided. The performance fees are based on a comparison of the Goldman roll strategy and the adaptive multiple-day roll strategy.

Table 7 – Intraday performance fees (adaptive single-day roll strategy v. adaptive multiple-day roll strategy)

Preference ($1 - \theta$)	Statistic							
	MN	SD	IQR	5%	10%	50%	90%	95%
(a) Corn Data: Early trading session (8.30am to 10.00am)								
0.1 (mean-focused)	0.26	0.26	0.21	0.62	0.50	0.24	0.04	-0.02
0.3	0.51	0.90	0.33	1.27	0.91	0.37	0.14	0.07
0.5	0.92	1.63	0.63	2.57	1.75	0.60	0.23	0.13
0.7	1.94	5.15	1.41	5.77	3.82	1.17	0.38	0.18
0.9 (variance-focused)	6.89	17.27	5.27	21.34	13.89	3.92	1.02	0.28
(b) Corn Data: Mid trading session (10.01am to 12.00pm)								
0.1	0.38	0.26	0.21	0.72	0.60	0.35	0.16	0.09
0.3	0.60	0.71	0.32	1.32	0.98	0.48	0.25	0.19
0.5	1.00	2.61	0.60	2.54	1.78	0.70	0.34	0.24
0.7	1.93	3.76	1.33	5.49	3.70	1.22	0.48	0.29
0.9	6.64	16.84	5.06	20.40	13.39	3.83	1.05	0.35
(c) Corn Data: Late trading session (12.01pm to 1.15pm)								
0.1	0.19	0.36	0.24	0.61	0.45	0.15	-0.06	-0.13
0.3	0.49	1.04	0.39	1.41	0.97	0.32	0.06	-0.01
0.5	1.05	2.48	0.77	3.07	2.05	0.63	0.18	0.07
0.7	2.31	6.45	1.72	6.88	4.58	1.33	0.38	0.17
0.9	8.72	22.09	6.63	26.61	17.54	4.88	1.32	0.55

Notes: This table contains information relating to the distribution of the performance fee δ obtained using the Bayesian approach applied to corn futures bid-ask spread data observed at various points within the trading day. Specifically, the mean (MN), standard deviation (SD), interquartile range (IR), and 5, 10, 50, 90, and 95 percentiles are provided. The performance fees are based on a comparison of the adaptive single-day roll strategy and the adaptive multiple-day roll strategy.

Table 8 – Performance fees (Goldman roll strategy robustness check)

Preference ($1 - \theta$)	Statistic							
	MN	SD	IQR	5%	10%	50%	90%	95%
(a) Coffee data: Goldman roll strategy v. adaptive single-day roll strategy								
0.1 (mean-focused)	0.02	0.12	0.09	0.15	0.12	0.03	−0.08	−0.14
0.3	−0.07	0.30	0.13	0.11	0.08	−0.02	−0.23	−0.36
0.5	−0.22	0.65	0.24	0.08	0.04	−0.11	−0.54	−0.84
0.7	−0.60	1.80	0.52	0.07	−0.01	−0.31	−1.29	−2.00
0.9 (variance-focused)	−2.43	6.29	2.00	0.08	−0.22	−1.33	−5.09	−7.88
(b) Coffee data: Goldman roll strategy v. adaptive multiple-day roll strategy								
0.1	−0.05	0.04	0.04	0.01	−0.01	−0.05	−0.09	−0.11
0.3	−0.06	0.11	0.05	0.01	−0.01	−0.05	−0.11	−0.14
0.5	−0.07	0.16	0.07	0.05	0.01	−0.06	−0.16	−0.22
0.7	−0.10	0.38	0.13	0.15	0.06	−0.08	−0.28	−0.43
0.9	−0.26	1.62	0.48	0.71	0.36	−0.16	−0.95	−1.51
(c) Crude oil data: Goldman roll strategy v. adaptive single-day roll strategy								
0.1	0.11	0.08	0.10	0.24	0.21	0.11	0.02	−0.01
0.3	0.01	0.08	0.11	0.14	0.11	0.01	−0.10	−0.13
0.5	−0.18	0.11	0.15	−0.01	−0.04	−0.17	−0.32	−0.37
0.7	−0.61	0.20	0.26	−0.31	−0.37	−0.59	−0.87	−0.97
0.9	−2.78	0.74	0.94	−1.73	−1.92	−2.70	−3.74	−4.10
(d) Crude oil data: Goldman roll strategy v. adaptive multiple-day roll strategy								
0.1	−0.01	0.04	0.06	0.06	0.05	−0.01	−0.06	−0.08
0.3	−0.02	0.05	0.06	0.05	0.03	−0.02	−0.08	−0.10
0.5	−0.05	0.05	0.07	0.03	0.01	−0.05	−0.12	−0.14
0.7	−0.12	0.08	0.11	0.01	−0.02	−0.12	−0.23	−0.26
0.9	−0.46	0.28	0.36	−0.03	−0.13	−0.45	−0.82	−0.94

Notes: This table contains information relating to the distribution of the performance fee δ obtained using the Bayesian approach applied to coffee and crude oil futures bid-ask spread data observed using various benchmark strategies. Specifically, the mean (MN), standard deviation (SD), interquartile range (IR), and 5, 10, 50, 90, and 95 percentiles are provided. The performance fees are based on a comparison of the Goldman roll strategy with adaptive single-day and multiple-day roll strategies.

Figure 1 – Performance measurement

This figure contains a graphical representation of two different ways of calculating the variance reductions (for a fixed mean level) available to users of the MVE roll strategy.

Figure 2 – Execution quality profiles (bid-ask spread data)

This figure contains the mean, autocorrelation (first-order), and variance of execution quality associated with all *s*-day and MVE roll strategies during the 24-day window prior to the FND. All panels provide information based on all grain futures observed over all trading periods.

Figure 3 – Execution quality profiles (illiquidity ratio data)

This figure contains the mean, autocorrelation (first-order), and variance of execution quality associated with all s -day and MVE roll strategies during the 24-day window prior to the FND. All panels provide information based on all grain futures observed over all trading periods.

Figure 4 – Execution quality profiles (intraday variation)

This figure contains the mean, autocorrelation (first-order), and variance of execution quality associated with all *s*-day and MVE roll strategies during the 24-day window prior to the FND. All panels provide information based on corn futures bid-ask spread data observed during various periods within each trading day.

Figure 5 – Execution quality profiles (coffee futures data)

This figure contains the mean, autocorrelation (first-order), and variance of execution quality associated with all *s*-day and MVE roll strategies during the 24-day window prior to the FND. All panels provide information based on coffee futures bid-ask spread and illiquidity ratio data.