ABSTRACT OF:

Catching up with the trend: A lagged cross-sectional momentum strategy for commodity futures

February 28, 2014

SUBMITTED FOR REVIEW TO THE

NATIONAL ASSOCIATION OF ACTIVE INVESTMENT MANAGERS

WAGNER AWARD 2014

By

Anshuman Mishra

The MathWorks, Inc.

Anshuman.Mishra@uncbusiness.net

ABSTRACT

Earlier research work has demonstrated the efficacy of absolute momentum and cross-sectional momentum-based trading strategies in the realm of commodity futures. Commodity markets are amenable to this form of trading since they exhibit sustained trends driven by economic and geopolitical factors. Groups of commodities tend to exhibit correlated behavior during such trends, though not all commodities in the group move together at the same time or at the same pace.

In this paper we propose a lagged cross-sectional momentum strategy to take advantage of this difference in trending behavior; our strategy demonstrates significant outperformance on a variety of investment performance metrics, compared to a benchmark portfolio of equally-weighted commodity futures in out-of-sample tests. Specifically, on out-of-sample testing from October 2004 to January 2014, compared to the benchmark numbers, our strategy boosts our monthly realized returns by nearly 50%, increases the monthly Sharpe ratio by 37% and the monthly Sortino ratio by 53%. This outperformance is accomplished with lower trading costs compared to the benchmark, since we only trade one commodity per group per month.

We begin our analysis by constructing a data set of 19 commodity futures contracts, categorized into 5 explicit groups – *energy*, *grains*, *soft* commodities, *metals* and *livestock*. Using only daily continuous front-month OHLC data, and maintaining strict data hygiene, we split our data set into a training set for exploratory analysis and strategy definition, a validation set for parameter optimization, and a test set for out-of-sample performance testing of our proposed strategy.

We proceed by validating the presence of significant correlations within explicitly identified commodity groups such as energy, metals, grains, livestock and soft commodities. Our thesis builds upon this correlated behavior by suggesting that within these groups, certain commodities react in a lagged fashion to group trends. We demonstrate the existence of this behavior on our explicitly. We posit that the reason behind this behavior lies in the realm of market psychology – for instance, hypothetically if gold and silver demonstrate a sustained trend, speculators might assume that platinum might behave the same in the near future, especially if it hasn't done so yet.

Next, we propose a cross-sectional momentum based trading strategy to take advantage of the outperformance of these laggard commodities, including a step-by-step walkthrough of the strategy's critical process of picking laggards. This strategy is parameterized by a single number: the size of the look-back window for momentum calculations. We pick one laggard per group, for a one-month holding period. We repeat this process each month.

Fully mapping out a trading strategy also involves design decisions such as estimating trading costs. For instance, since we compute momentum at the end of the month, from monthly close to monthly close, but trade on the next month's opening prices, we need to account for slippage costs incurred during this gap. We demonstrate one way of estimating this cost, one that can be applied to other trading frequencies.

Going beyond standard metrics such as Sortino ratio to evaluate strategy performance, we realize that in order to choose an optimal parameter, trader preferences and risk appetites need to be accounted for. Some traders can stomach larger mark-to-market losses in the form of drawdowns, while others want tighter risk controls. These qualitative preferences can be mapped into comparative level and sign preferences for moments of the returns distribution generated by any trading strategy being considered.

Our validation testing of the strategy involves designing a custom objective function that includes skewness and kurtosis preferences for determining the optimal parameter. We stress the fact that objective function design is dependent upon trader preferences; there is no “one size fits all” measure.

We perform out-of-sample testing of our proposed strategy on data that was realized historically from October 2004 to January 2014. However, this creates the possibility of an inaccurate out-of-sample test, because it is unlikely that traders will trade the strategy continuously during this period. Traders are more likely to stop trading the strategy during periods of underperformance. Accordingly, we account for this behavior by performing a Monte Carlo simulation of 5000 random paths of different durations (ranging from 1 month to years) within our out-of-sample data, and computing relevant statistics. This form of randomized start and stop testing of our strategy is more likely to closely match typical trader behavior.

We tabulate month-by-month comparative returns against the benchmark portfolio, and conclude by drawing out possible extensions of our work into other asset classes.

KEYWORDS: momentum trading, trend following, cross-sectional momentum, commodity futures

Catching up with the trend: A lagged cross-sectional momentum strategy for commodity futures

February 28, 2014

SUBMITTED FOR REVIEW TO THE

NATIONAL ASSOCIATION OF ACTIVE INVESTMENT MANAGERS

WAGNER AWARD 2014

By

Anshuman Mishra

The MathWorks, Inc.

Anshuman.Mishra@uncbusiness.net

1. Introduction

Momentum is a well-established anomaly in the financial literature as well as in professional practice [MOS2011]; there exists a significant amount of research to indicate the presence of momentum across asset classes (equities, currencies, commodities [MOS2012]), in a variety of market regimes [STI2013]. Multiple explanations have been posited to explain the persistence of the anomaly, with behavioral models of trading agents being the preferred approach [JEG1999].

Exchange-based commodity markets are a popular avenue for the testing and deployment of momentum-based trading strategies. Closely related trend-following funds that trade commodity futures have demonstrated persistent market-beating returns in this arena, and are an especially good hedge during periods of crises and market turmoil [HUT2014].

Research into momentum strategies for commodity markets exhibits two traits:

(1) the inclusion of exchange-traded instruments such as equity index futures or interest rate futures that are not, strictly speaking, commodities, and

(2) Treating all commodities equally though some tend to move together, in the sense of trending up or down

In this paper, we make the claim that certain commodities trend together in groups. Further, our next claim is that those commodities within these groups that lag in their returns as compared to the overall group tend to “catch up” by exhibiting higher returns than the rest of the group after a given period of time. We proceed to combine these claims into a trading strategy that consistently outperforms an equally-weighted buy-and-hold strategy on measures such as the Sortino ratio, while lowering maximum drawdown and downside volatility.

The rest of this paper is organized as follows. In section 2, we describe our data collection, cleansing and analytics setup. In section 3, we analyze our claim that certain commodities that lie in defined groups tend to move together. In section 4, we look at multiple ways of defining momentum and state our preferred approach. In section 5, we study the behavior of laggards within commodity groups and demonstrate that their performance is exploitable. In section 6, we define our trading strategy and perform validation testing to pick the appropriate strategy parameters. In section 7, we analyze the performance of the strategy in out-of-sample tests, and highlight particular aspects of the trading strategy. In section 8, we conclude by summarizing our technique and proposing certain enhancements to the trading strategy.

2. Data Collection and Analytics Setup

We collect continuous daily front-month futures OHLCV data from Quandl.com for the 19 commodities listed in Table 1 and bucket them into groups. The retrieved data is then synchronized to dates in common across all commodities. In effect we have daily data for these 19 commodities from 17th August 1990 to 19th February 2014.

The next step in our data pre-processing is to roll up the obtained daily data into monthly OHLC prices, and place the data into the groups defined in Table 1.

Finally, we split the data into a training set (30% of the data), a validation set (30%), and a test set (40%). The training data (Aug 1990 – Jul 1997) is used to set the case for claims we make regarding lagged trending behavior within commodity groups. The validation data (Aug 1997 – Jul 2004) is used to pick optimal strategy parameter ranges, and the test set (Aug 2004 – Feb 2014) is used for out-of-sample strategy testing. As per established practice [MOS2011], we assume our futures positions to be fully collateralized.

Group	Commodity (SYMBOL)	Number
Energy	WTI Crude Oil (CL), Brent Crude Oil (B), Natural Gas (NG), Heating Oil (HO)	4
Grains	Corn (C), Wheat (W), Soybean (S), Soybean Oil (BO)	4
Softs	Sugar (SB), Coffee (KC), Cocoa (CC), Cotton (CT)	4
Metals	Gold (GC), Silver (SI), Copper (HG), Platinum (PL), Palladium (PA)	5
Livestock	Live Cattle (LC), Lean Hogs (LN)	2

TABLE 1: Selected commodity contracts

3. How strongly do these commodities move together?

Though we have identified certain commodities as belonging to particular groups, there is no reason to expect that these commodity groups will exhibit strong internal co-movements in the evolution of their price series. In order to generate some visual intuition around such co-movements, we plot the price evolution of each commodity group in charts and then analyze their internal correlations.

Figure 1(a)-(e) is a visual comparison of the commodities' price evolutions within the training dataset, highlighting certain aspects of their behavior:

(1) Some groups exhibit little or no internal co-movement. This stands to reason since, beyond the vagaries of weather, there are widely different economic forces driving the prices of, say, coffee and cotton.

(2) Some commodities (*Energy*, *Metals*) tend to move together for certain sub-periods, with prominent exceptions (natural gas NG, palladium PA) that do not follow the trends set by the rest of the group.

Correlation analysis. We perform a Spearman rank correlation analysis [SEW2011] within and across groups. The correlations are computed on pairs of commodities, both within and across groups, on daily returns for the training period, and the average group correlations are computed and reported in Table 2. The Spearman statistic is preferable to the commonly-used Pearson statistic as the former is non-parametric and can be used even for nonlinear forms of monotonic dependence.

	Energy	Grains	Softs	Metals	Livestock
Energy	**0.4619**				
Grains	0.0164	**0.4775**			
Softs	0.0049	0.0690	0.0659		
Metals	0.0547	0.0619	0.0495	**0.3659**	
Livestock	0.0108	0.0310	0.0230	0.0275	**0.2004**

TABLE 2: Average Group Correlations

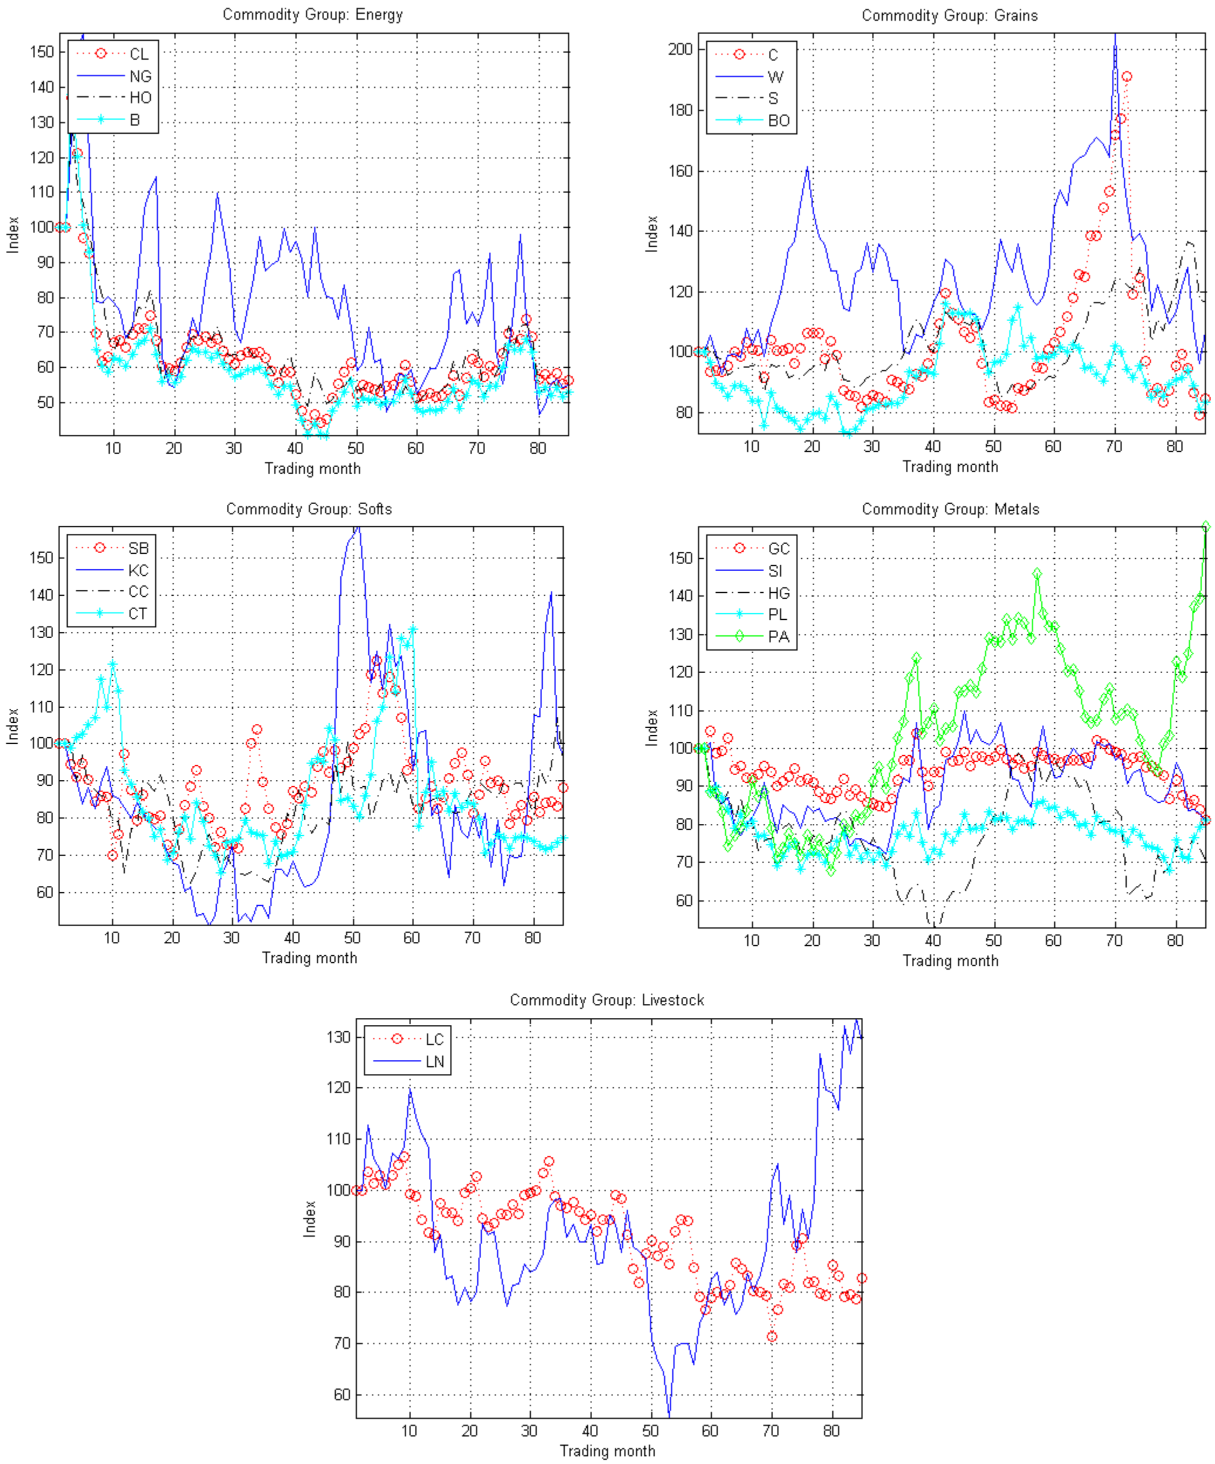

<u>FIGURE 1</u>: Commodity price evolution across groups
[L-R, Top to Bottom: *Energy*, *Grains*, *Softs*, *Metals*, *Livestock*]

The results of the correlation analysis indicate the following:

(1) Groups such as *Energy* and *Grains* display high internal correlation.

(2) *Metals* and *Livestock* display a moderate amount of internal correlation.

(3) *Soft* commodities exhibit negligible internal correlation.

(4) Correlation across groups is negligible or nearly zero.

4. Measuring momentum

There are two popular ways to measure momentum for the purposes of trading: absolute momentum, defined for a single time-series of asset data without reference to any other series, and relative momentum, defined on a cross-sectional basis that ranks returns for multiple assets at a single point in time.

Absolute momentum. The simplest way to define absolute momentum [ANT2013] of an asset is by measuring its return over a given look-back period *k*, starting at time *t*.

$$ABSMOM = log\frac{Close_{t-1}}{Close_{t-k-1}} \quad (1)$$

As per standard techniques, we skip 1 month when computing the metric in order to guard against look-ahead bias.

Other ways of measuring absolute momentum include measuring the instantaneous slope of the simple moving average, or the difference between a fast moving average and a slow moving average.

Relative momentum. In the interests of space, we focus on using (1) as our absolute momentum measure; we proceed to define relative momentum on the basis of this measure as well. Given a group of commodities G = {C_1, C_2, C_3, …, C_n} and a look-back window of size *k*, we compute $ABSMOM_i$ for all C_i, and rank the absolute momentum numbers from lowest to highest.

5. Do some commodities lag others within groups?

In this section of the paper, we use the cross-sectional momentum metric defined previously to measure the subsequent performance of laggards, i.e. commodities that have the lowest relative momentum within their group.

Our procedure is setup as follows:

1. Our data set is monthly training data from August 1990 to July 1997, for the 19 commodities categorized into 5 groups as per Table 1.
2. For a given look-back window *k* (say 6 months), we walk through the test data month by month starting at *k+1*, computing the absolute momentum

numbers for each commodity in each group. As before, we add 1 month to guard against look-ahead bias.

3. For each group, we identify the laggard commodity and note its performance in the current month.

Table 3(a)-(c) demonstrates an application of the above procedure, using a look-back window of 6 months, for the group of grain commodities.

The procedure outlined above uncovers significant outperformance due to the lagged momentum phenomenon. We believe the reasons that lead to the existence of this phenomenon lie in the realm of behavioral finance and market psychology. Our operating thesis is that market participants believe certain commodities (within groups such as *grains* or *energy*) move together; in the presence of strong upside momentum for a particular commodity group, speculators desiring to participate choose the most closely-related commodity that is still relatively cheap – this turns out to be the laggard commodity.

We run the procedure on the training data, constructing a returns series for each group, consisting of the returns of the laggards held for a month. We take into account the additional cost of active trading as compared to the equally-weighted

(a) CLOSING PRICES

	Corn [C]	Wheat [W]	Soybean [S]	Soybean Oil [BO]
Jul-96	354.25	440.00	765.75	24.25
Aug-96	370.50	448.50	814.50	25.33
Sep-96	296.75	436.00	758.00	23.86
Oct-96	266.00	371.25	667.75	22.59
Nov-96	272.00	400.50	715.25	23.23
Dec-96	258.25	381.25	690.50	22.71
Jan-97	270.25	359.75	738.25	23.76
Feb-97	296.75	373.00	791.00	24.22
Mar-97	308.25	401.75	855.00	24.42
Apr-97	295.00	423.00	889.50	25.06
May-97	270.75	360.50	880.50	23.78
Jun-97	248.00	323.75	771.00	21.79
Jul-97	265.50	362.00	768.00	22.35

(b) RETURNS

	C	W	S	BO
Jul-96	-37.66%	-8.39%	-1.27%	-2.93%
Aug-96	4.49%	1.91%	6.17%	4.36%
Sep-96	-22.20%	-2.83%	-7.19%	-5.98%
Oct-96	-10.94%	-16.08%	-12.68%	-5.47%
Nov-96	2.23%	7.58%	6.87%	2.79%
Dec-96	-5.19%	-4.93%	-3.52%	-2.26%
Jan-97	4.54%	-5.80%	6.69%	4.52%
Feb-97	9.35%	3.62%	6.90%	1.92%
Mar-97	3.80%	7.43%	7.78%	0.82%
Apr-97	-4.39%	5.15%	3.96%	2.59%
May-97	-8.58%	-15.99%	-1.02%	-5.24%
Jun-97	-8.78%	-10.75%	-13.28%	-8.74%
Jul-97	6.82%	11.17%	-0.39%	2.54%

(c) COMPUTING MOMENTUM

	C	W	S	BO
Dec-96	-5.19%	-4.93%	-3.52%	-2.26%
Jan-97	4.54%	-5.80%	6.69%	4.52%
Feb-97	9.35%	3.62%	6.90%	1.92%
Mar-97	3.80%	7.43%	7.78%	0.82%
Apr-97	-4.39%	5.15%	3.96%	2.59%
May-97	-8.58%	-15.99%	-1.02%	-5.24%
ABSOLUTE MOMENTUM	-0.46%	-10.52%	20.79%	2.34%
RELATIVE RANK	3	4	1	2

SKIP MONTH	C	W	S	BO
Jun-97	-8.78%	-10.75%	-13.28%	-8.74%

LAGGARD Wheat

JULY 1997 RETURNS	C	**W**	S	BO
Jul-97	6.82%	**11.17%**	-0.39%	2.54%

WALKTHROUGH
Look-back window size: 6 months

From (a), we compute monthly close-to-close returns as shown in (b).
Assuming we are at the beginning of July 1997; looking back, we skip 1 month (June 1997), and, in (c), we compute the total close-to-close returns from November 1996 to May 1997, by summing up the returns from December 1996 to May 1997. We then assign relative ranks to each of the absolute momentum numbers on a comparative basis. We thus identify the laggard commodity, i.e. the commodity that has the least momentum, as **wheat.**
Going long on the laggard commodity (Wheat) in July generates the maximum returns for that month, in the group of grain commodities.

<u>TABLE 3</u>: An example of identifying lagging commodities portfolio by subtracting 25 basis points from each month's laggard returns. Next, for each group, we prepare a comparative plot of the P&L curve of the laggards' returns series against the P&L curve of an equally-weighted buy-and-hold benchmark portfolio that invests in every commodity in that group. Finally, we compute a P&L curve that is equally-weighted across each of the five groups.

Figure 2 presents the results of our procedure; each of the six charts compares the P&L curves generated by the laggards with that of the equally-weighted portfolio of the commodity group. Note that the last chart computes P&L curves that are equally-weighted across all groups, for both the laggards' returns series as well as that of the benchmark.

We note that:

(1) Our strategy of trading the momentum laggard outperformed the corresponding benchmark for *Energy*, *Softs*, and *Livestock*, and underperformed in the case of *Grains* and *Metals*.

(2) Weighting our strategy equally across all groups outperformed the benchmark.

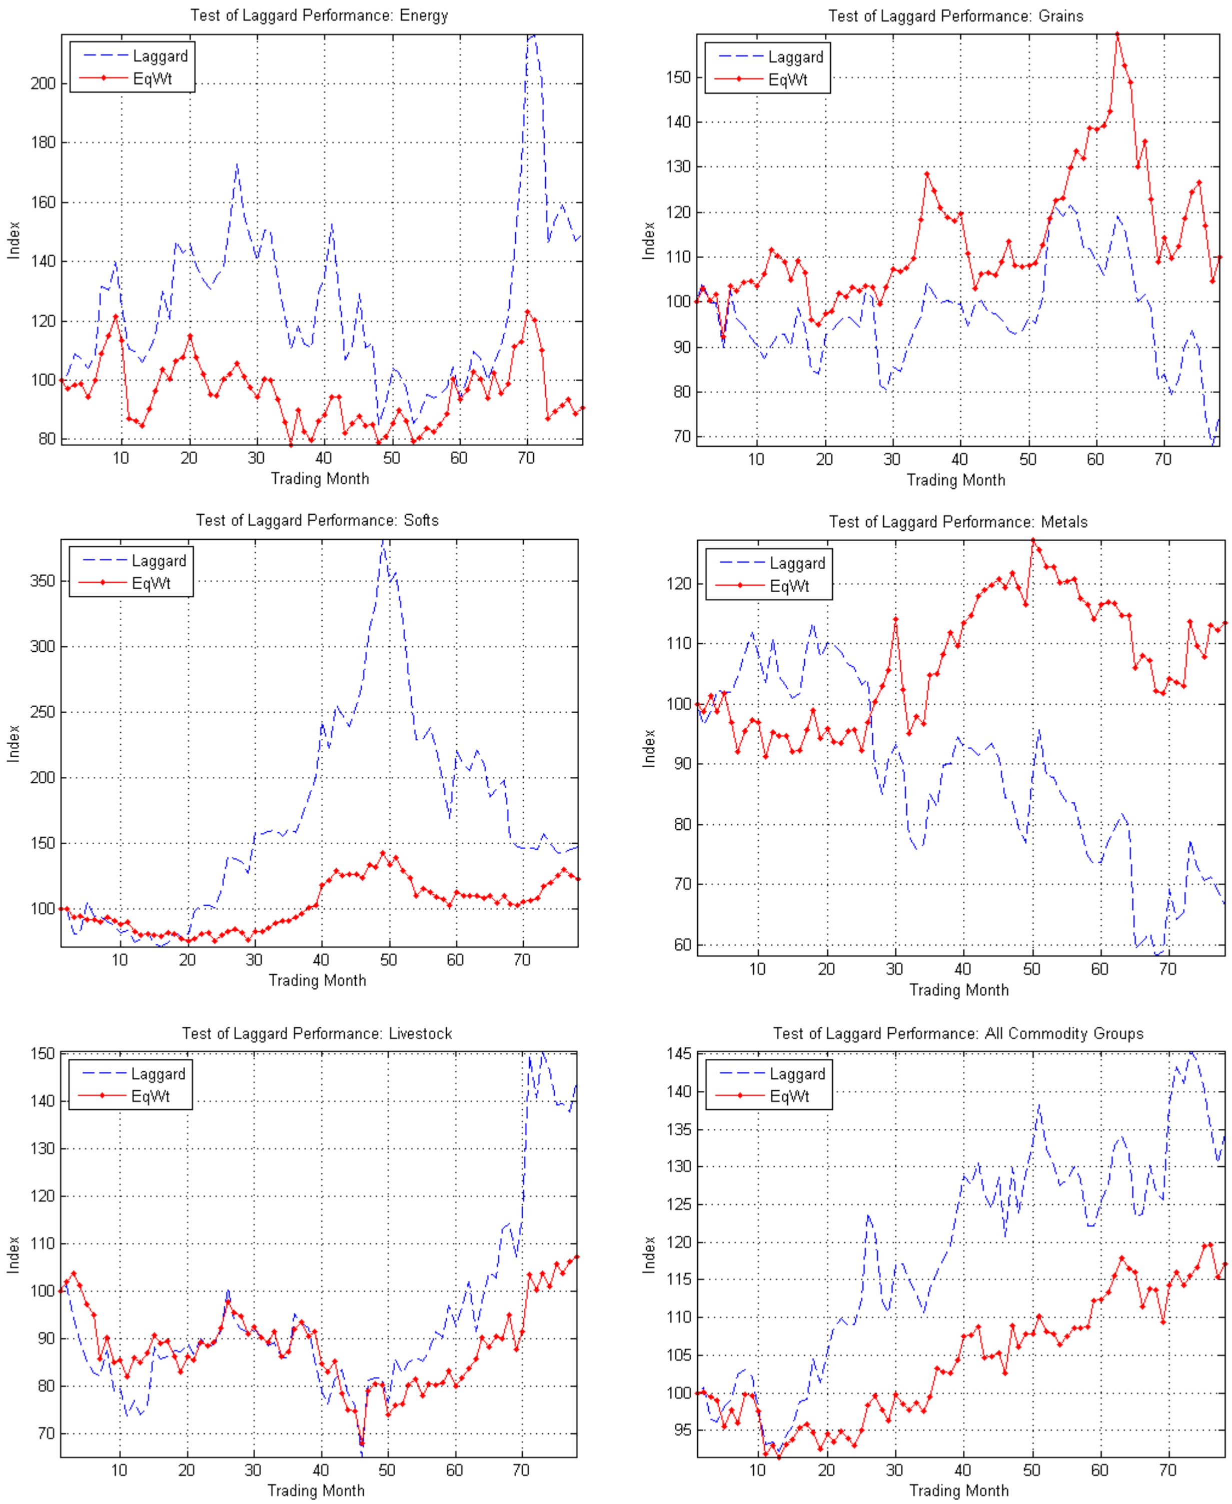

FIGURE 2: Comparative P&L curves for all groups, Aug 1990-Jul 1997

[Left to Right, Top to Bottom: *Energy, Grains, Softs, Metals, Livestock, All Groups.*

Dashed lines indicate our strategy; dotted lines are for equal-weight benchmark]

Performance metrics. In order to quantify our intuition, we compute relevant performance metrics (see Table 4). An interesting metric to note is the *Return Source %*; this measures the relative frequency with which each commodity being identified as a laggard within the group. Note that we measure the Sortino ratio using a minimum acceptable return of zero, and set the risk-free rate to zero for Sharpe ratio calculations. Both ratios are calculated on a monthly basis.

Examining the *Return Source %* metric in Table 4 leads us to conclude that group-wise strategy returns are not dominated by any single commodity.

	Energy	Grains	Softs	Metals	Livestock
Avg Mthly Return	1.022% (0.118%)	-0.169% (0.243%)	0.997% (0.371%)	-0.323% (0.229%)	0.705% (0.188%)
Std. Dev. Of Returns	10.021% (6.864%)	6.240% (4.872%)	10.179% (4.585%)	6.348% (3.603%)	6.957% (4.500%)
Maximum Drawdown	51% (35%)	44% (34%)	63% (28%)	49% (20%)	36% (35%)
Sharpe Ratio	0.103 (0.017)	-0.027 (0.050)	0.099 (0.081)	-0.051 (0.064)	0.102 (0.042)
Sortino Ratio	0.156 (0.023)	-0.037 (0.069)	0.168 (0.127)	-0.070 (0.097)	0.194 (0.063)
Return Source %	CL:23%,NG:38% HO:18%,B:21%	C:26%, W:30% S:10%, BO:34%	SB:22%, KC:30% CC:16%, CT:32%	GC:17%, SI:19% HG:36%, PL:6% PA: 21%	LC:57% LN:43%

TABLE 4: Grouped performance metrics of laggard commodity returns

[Look-back window of 6 months; numbers in parentheses are equivalent metrics for equal-weighted benchmark for that group]

Equal weighting across groups. Though trading the momentum laggard outperforms the benchmark in three groups (higher Sortino ratios), we still face a larger drawdown in most cases. Also, since we have no way of knowing ahead of time which commodity group our strategy will outperform in, it makes sense to trade our strategy equally across all groups, i.e. allocate 20% weights to momentum laggards in each group and measure the performance. The benchmark in this case will be an equally-weighted portfolio across all 19 commodities.

As before, we test the strategy outlined above on monthly training data; the results are collected in Table 5. We use the same momentum look-back period of 6 months.

	Avg Mthly Return	Std. Dev. of Returns	Maximum Drawdown	Mthly Sharpe Ratio	Mthly Sortino Ratio
Strategy	0.446%	3.506%	12%	0.128	0.208
Benchmark	0.230%	2.177%	9%	0.106	0.161
Ratio	1.93	1.61	1.33	1.207	1.292

TABLE 5: Equally-weighted strategy vs. benchmark

The results indicate our strategy beats the benchmark by a factor of 2 in terms of the monthly return, compensating for the increase in standard deviation.

However, most of the volatility appears to be on the upside, as the Sortino ratio is boosted by nearly 30%, and the Sharpe ratio is increased by 20%.

In the next section, we fully lay out the trading strategy, measure its performance metrics on validation data that is independent of the training data, and pick the best value for the size of the look-back window. We then perform out-of-sample tests on the trading strategy and examine the evidence for outperformance.

6. Defining the trading strategy

We propose the following long-only trading strategy to exploit the opportunities uncovered by analyzing the relative momentum of laggard commodities:

1. Collect monthly closing data for the 19 commodities defined earlier.
2. Select a look-back window size, say *k* months. Assume round-trip trading costs of 25 basis points per trade.
3. At the beginning of every month:
    a. For each group G_i (i = 1..5)
        i. Skip the previous month's data.
        ii. Measure the absolute momentum over the lookback window for each commodity in the group.

iii. Rank the momentum measures and identify the laggard commodity L_i.

b. Allocate an equal amount of capital *F* across each group G_i.

c. Buy every commodity L_i (i = 1..5) on a fully-collateralized basis at the monthly open.

4. Hold all L_i until the end of the month and sell at the monthly close.
5. Repeat this procedure every month.

We propose a corresponding benchmark for this strategy, which is simply an equally-weighted portfolio of all 19 commodities, rebalanced monthly on a trading friction-free basis.

Accounting for trading costs. Two components that we consider in our analysis are transaction costs and slippage costs. In our strategy, while we measure returns on a close to close basis, we execute our trades from the monthly open to the monthly close. To account for this discrepancy, we analyze the distribution of average slippage costs (month close to next month open) across all commodities.

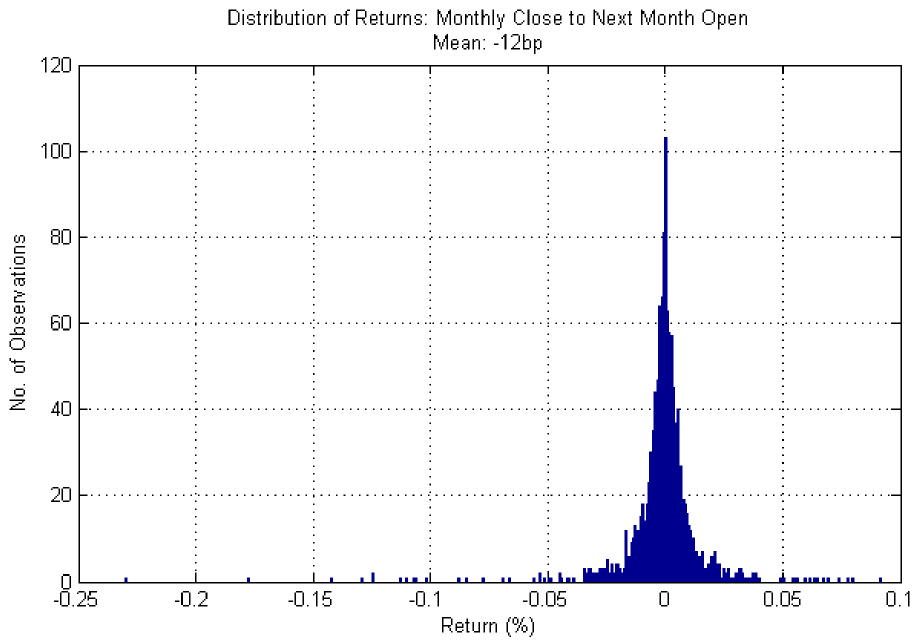

FIGURE 3: Analyzing trading slippage costs

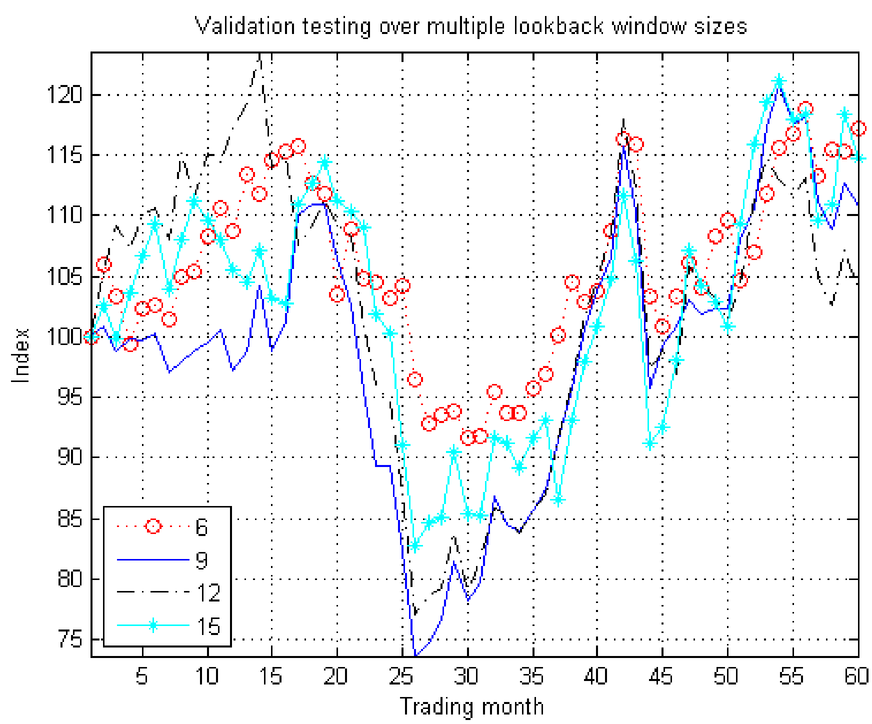

FIGURE 4: Validation testing over multiple look-back window sizes

Figure 3 displays a histogram of the slippage return distribution; the mean slippage is 12 basis pts, though the returns are tightly clustered around zero.

We assume transaction costs to be of the same order of magnitude, and hence double our slippage estimate to arrive at the combined trading costs. Our final cost per trade is thus assumed to be 25 basis points.

Validation testing and parameter optimization. We test the strategy with a range of look-back window sizes (see Figure 4), on the validation data set defined in section 2, consisting of data for all commodities from August 1997 to July 2004. An important point to note is that there is no overlap between the training data used to construct the trading strategy, and the validation data used to perform strategy optimization.

Though we have only one parameter to optimize for our strategy, namely the size of the look-back window, it is best to figure out what we should be optimizing for. This boils down to a set of preferences that we need to quantify and combine into a single function for optimization purposes.

Needless to say, our first preference is to maximize the average monthly returns generated by the portfolio. However, going with this preference alone subjects us to the possibility of facing large drawdowns in our capital. Therefore, our second preference is to minimize downside risk by minimizing the drawdowns exhibited on the validation test data. We estimate the mean realized drawdown over rolling 12-month periods, and divide it by 12 to arrive at a monthly number.

Though not exactly congruent, optimizing for these two preferences can be interpreted as selecting the parameter that yields the maximum Sortino ratio. However, it pays to further analyze the situation – investment management performance is impacted to a great degree by the emergence of outliers. Exposure to a trading strategy that has a positively-skewed returns distribution is beneficial since emergent outliers are more likely to be positive returns rather than negative returns. Accordingly, our third preference is to maximize positive skewness in our strategy's returns distribution.

Finally, after the financial crisis, a prime concern is to hedge against fat tail risk, i.e. the chance that there is a greater than normal probability of experiencing returns that lie at extremes to the mean. This is important because, if we have a trading strategy that has a stable positive expected return, ideally we would like it to continue generating returns equal to or close to it. Fat tail risk is embedded in the kurtosis of the strategy's returns distribution. This leads us to our fourth preference: to minimize kurtosis.

Of course, changes in market behavior over time will lead to different optimal parameters. We can partially hedge against this outcome by selecting those parameters that are surrounded by other parameters that generate close-to-optimal performance, or, in the parlance of optimization terminology, we select parameters that lie on a *plateau* of the solution surface, not a *sharp peak*.

To recap, our preferences are to:

(1) Maximize returns

(2) Minimize drawdowns

(3) Maximize positive skew

(4) Minimize kurtosis

Every trader will have a different set of preferences. One can model this by assigning relative weights to each of the preferences in our objective function; we will skip this step in the interests of a cleaner exposition.

The objective function below is evaluated by running the trading strategy for each look-back window *k*, ranging from 1 to 24, and measuring mean return, maximum drawdown, skewness and kurtosis.

$$StrategyObjective = \frac{Mean\ Return}{Mean\ Realized\ Drawdown} * \frac{Skewness}{Kurtosis} \quad (2)$$

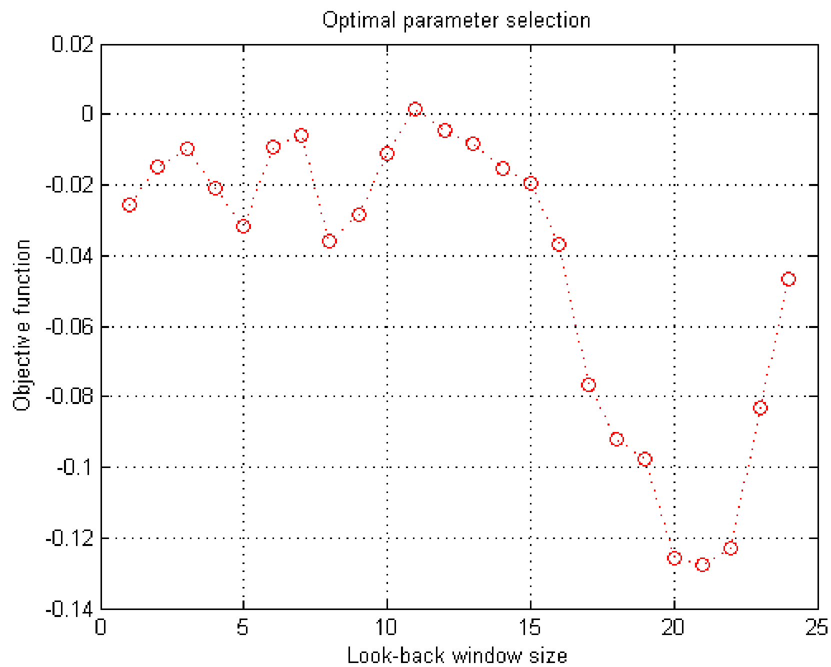

<u>FIGURE 5</u>: Selecting the optimal look-back window size

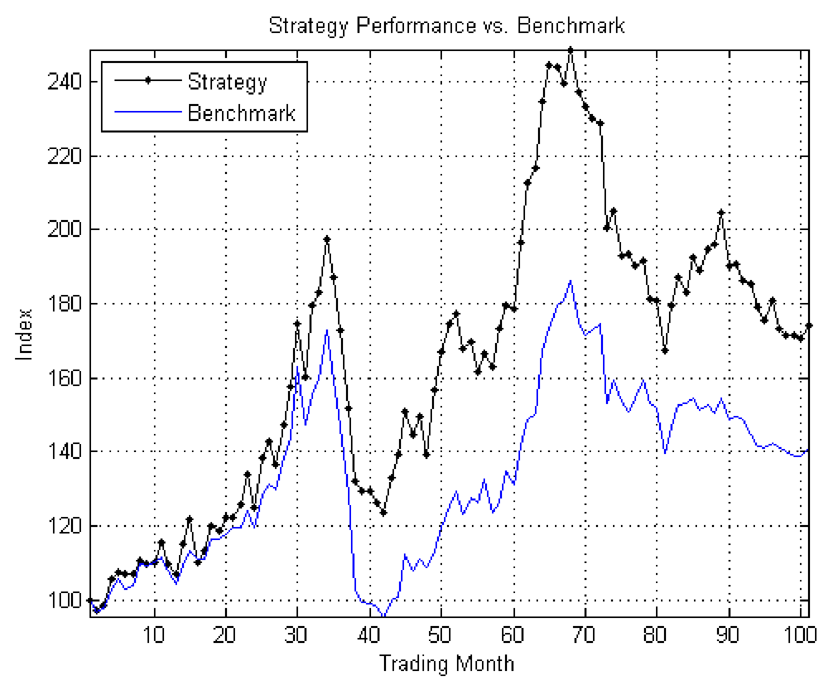

<u>FIGURE 6</u>: Out-of-sample strategy performance vs. benchmark

Figure 5 plots the strategy objective function for all look-back window sizes in our range.

We select our optimal look-back window size *k* to be 11 months, as it has the highest value for the objective function and is moreover located in a region that is stable. Note that strategy performance as per our stated preferences starts decaying drastically beyond a look-back window size of 15 months.

7. Testing the trading strategy

Out-of-sample testing. Now that we have selected our optimal parameter, we test our trading strategy on out-of-sample data. This consists of monthly data for all 19 commodities, from October 2004 to January 2014. As before, our benchmark strategy is an equally-weighted portfolio of all commodities.

Figure 6 reveals the outperformance of our strategy compared to the benchmark. Table 6 tabulates the comparative performance of our strategy. Compared to the benchmark, our strategy has an increased average monthly return, a lower standard deviation, a lower maximum drawdown, and a boosted Sortino ratio.

In Table 7, we capture the excess returns generated by our strategy compared to the benchmark portfolio, and tabulate the overall yearly outperformance. We

	Avg Mthly Return	Std. Dev. of Returns	Maximum Drawdown	Mthly Sharpe Ratio	Mthly Sortino Ratio
Strategy	0.699%	3.588%	38%	0.130	0.202
Benchmark	0.471%	5.004%	45%	0.095	0.132
Ratio	1.48	0.72	0.84	1.37	1.53

TABLE 6: Out-of-sample test of strategy vs. benchmark

	Year	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
2005	3.04%										0.53%	0.68%	1.81%
2006	-2.34%	-1.13%	1.99%	-1.05%	-2.25%	-0.57%	0.21%	3.61%	-2.22%	1.10%	2.93%	1.76%	**-7.39%**
2007	5.03%	2.89%	1.05%	-0.91%	1.87%	-1.67%	2.94%	2.87%	-3.37%	3.11%	1.21%	-3.43%	1.56%
2008	17.23%	2.80%	-2.25%	1.35%	**6.78%**	-1.21%	-0.05%	2.30%	0.29%	0.18%	**7.73%**	0.72%	0.64%
2009	6.53%	-1.83%	0.34%	3.46%	3.64%	-2.97%	-0.36%	0.61%	**-4.83%**	**8.14%**	0.95%	-0.13%	-1.89%
2010	2.01%	-0.22%	-2.50%	-3.54%	-2.31%	**4.65%**	**4.39%**	-3.32%	2.08%	1.81%	3.80%	0.48%	-3.02%
2011	-9.28%	0.64%	-3.57%	-2.70%	0.96%	1.42%	0.28%	-2.20%	-1.70%	0.03%	-1.68%	-2.30%	1.95%
2012	5.50%	-3.85%	-2.66%	-1.38%	0.31%	0.81%	1.90%	0.25%	-2.17%	**4.20%**	0.04%	2.33%	1.94%
2013	-7.10%	1.83%	-3.38%	0.01%	-1.90%	2.05%	-0.96%	-1.67%	1.93%	-3.24%	-0.51%	1.08%	-0.59%

TABLE 7: Month-by-month outperformance statistics for out-of-sample test

observe the excellent outperformance during turbulent market periods such as the financial crisis of 2008, lending credence to our strategy being positioned as a diversification vehicle.

Randomized start and stop testing. Computing the out-of-sample performance on our data set assumes necessarily that the trader will start trading at the beginning of the data set and stop trading at the end. This does not take into account the fact that implementations of this strategy might start trading at some random point in time and stop at some other random point in time.

We analyze our strategy on a randomized start and stop basis by generating 5000 pairs of points in time between October 2004 and January 2014 and computing corresponding simulated P&L paths for each pair; we then calculate the average monthly realized returns over each of these paths. Figure 7 is a histogram of our simulated returns. Table 8 lists relevant moment statistics.

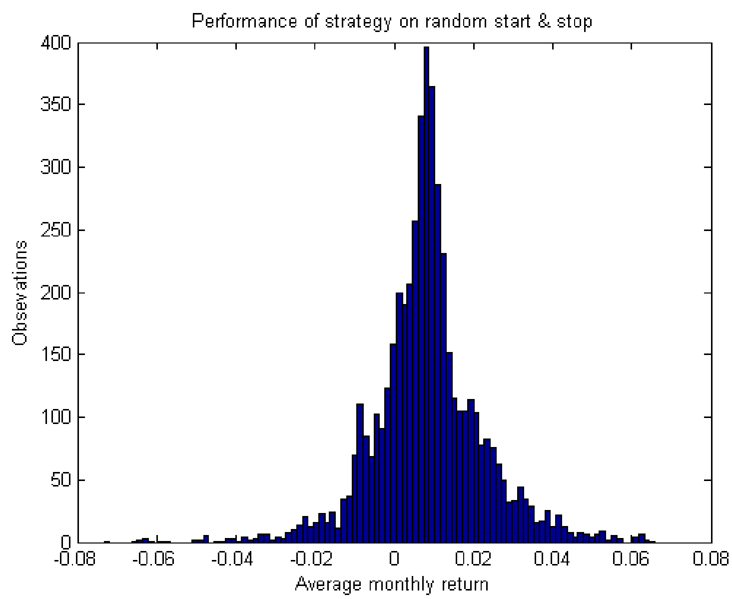

FIGURE 7: Randomized start and stop testing

	Median Mthly Return	Median Std. Dev.	Median Skewness	Median Kurtosis	Mthly Sharpe Ratio
Strategy	0.88%	5.68%	-0.20	2.45	0.155
Benchmark	0.59%	5.30%	-0.69	3.77	0.111

TABLE 8: Moment statistics for randomized start and stop testing

In practice, we expect our trading strategy performance metrics to be similar to the statistics generated from randomized start and stop testing. Our strategy has a nearly 50% higher monthly return, a higher skewness, and a lower kurtosis, at the expense of a higher standard deviation. On the implementation front, our strategy induces less trading friction (since only one commodity is traded per group leading to less transactions per period); a constant holding period of 1 month ensures we can comfortably switch the strategy on or off without excess operational overhead.

8. Conclusion

Starting from the popular momentum technique for trading strategies, this paper uncovers a lagged trending effect within commodities that generates systematic outperformance. We proceed to design a trading strategy that takes advantage of this effect to produce superior performance compared to a conventional equally weighted buy-and-hold commodity portfolio, and demonstrate its efficacy on out-of-sample market data, generating lower drawdowns, higher returns and higher Sortino ratios.

We believe this strategy will prove to be a useful building block when constructing commodity portfolios, considering its diversification benefits.

Extensions. The core of this strategy lies in explicitly identifying groups within commodity futures. A straightforward extension to this approach can be implemented by including other groups into the mix such as bond futures, equity index futures, or currency futures. A more involved extension would be to use this technique on baskets of actively-traded equities; in this case, since our datasets consist of greater number of symbols, statistical learning-based automated clustering methods can be used to take the manual guesswork out of group identification. Our preliminary results using multidimensional scaling techniques have been encouraging and further research in this direction points to potential rewards.

REFERENCES

[ANT2013] Antonacci, G., "Absolute momentum: A simple rule-based strategy and universal trend-following overlay", *Available at SSRN: http://ssrn.com/abstract=2244633,* 2013

[HUT2014] Hutchison, M. C., O'Brien, J. J., "Is this time different? Trend following and financial crises", *Available at SSRN: http://ssrn.com/abstract=2375733*, 2014

[JEG1999] Jegadeesh, N., Titman, S., “Profitability of momentum strategies: an evaluation of alternative explanations”, *National Bureau of Economic Research Working Paper 7159*, 1999

[MOS2011] Moskowitz, T., Ooi, Y. H, Pedersen, L., “Time Series Momentum”, *Journal of Financial Economics*, 2011

[MOS2012] Moskowitz, T., Asness, C., Pederson, “Value and Momentum Everywhere”, *Journal of Finance*, 2012

[SEW2011] Sewell, M., “Characterization of Financial Time Series”, UCL Department of Computer Science Research Note RN/11/01, 2011

[STI2013] Stivers, C., Sun, L., “Market Cycles and the Performance of Relative Strength Strategies”, *Financial Management*, 2013