



AgEcon SEARCH
RESEARCH IN AGRICULTURAL & APPLIED ECONOMICS

The World's Largest Open Access Agricultural & Applied Economics Digital Library

This document is discoverable and free to researchers across the globe due to the work of AgEcon Search.

Help ensure our sustainability.

Give to AgEcon Search

AgEcon Search

<http://ageconsearch.umn.edu>

aesearch@umn.edu

*Papers downloaded from **AgEcon Search** may be used for non-commercial purposes and personal study only. No other use, including posting to another Internet site, is permitted without permission from the copyright owner (not AgEcon Search), or as allowed under the provisions of Fair Use, U.S. Copyright Act, Title 17 U.S.C.*

**Technical Analysis in Commodity Markets:
Risk, Returns, and Value**

Matthew C. Roberts*

*Paper presented at the NCR-134 Conference on Applied Commodity Price Analysis,
Forecasting, and Market Risk Management.
St. Louis, Missouri, April 21-22, 2003*

Copyright 2003 by Matthew C. Roberts. All rights reserved. Readers may make verbatim copies of this document for non-commercial purposes by any means, provided that this copyright notice appears on all such copies.

*Assistant Professor, Department of Agricultural, Environmental, and Development Economics, The Ohio State University. 320 Ag. Admin. Bldg., 2120 Fyffe Rd., Columbus, Ohio 43210-1067. Phone: (614) 688-8686. Fax: (614) 292-0078. email: roberts.628@osu.edu. This manuscript dated May 19, 2003

Technical Analysis in Commodity Markets: Risk, Returns, and Value

Abstract

Although there is little academic research that supports the usefulness of technical analysis, its use remains widespread in commodity markets. Much prior research into technical analysis suffered from data-snooping biases. Using genetic programming, *ex ante* optimal technical trading strategies are identified. Because they are mechanically generated from simple arithmetic operators, they are free of the data-snooping bias common in technical analysis research. These rules are clearly capable of forecasting periods of high and low volatility, but rules generated for corn and soybeans cannot consistently generate profits in the presence of transactions costs. Rules generated for wheat futures produce profits that are weakly significant, both statistically and economically.

Keywords: Technical Analysis, Genetic Algorithms, Commodity Markets, Futures Markets

Technical analysis is a broad collection of methods and strategies which attempt to forecast future prices on the basis of past prices or other observable market statistics, such as volume or open interest. Based on this definition, technical analysis conflicts with weak-form market efficiency, under which “efficiency with respect to an information set . . . implies that it is impossible to make economic profits by trading on the basis of [that information set],” (Malkiel) and the information set consists of precisely the information which technical analysis purports to exploit.

Academia maintains a generally negative view of technical analysis, perhaps best typified by Malkiel, “Obviously, I am biased against the chartist. This is not only a personal predilection, but a professional one as well. Technical analysis is anathema to the academic world.” Although there are some that are more charitable toward technical analysis, Campbell, Lo, and MacKinlay suggest that “perhaps some of the prejudice against technical analysis can be attributed to semantics.” Nevertheless, the study of technical analysis has a long history in academia, with mixed results.

Early studies, such as Alexander and Fama and Blume identified and tested simple technical strategies using equity index data and found that although they may have some predictive power, they were unable to consistently generate positive profits. Over the succeeding decades, similar conclusions were reached by many researchers, especially when transactions costs were included in the analysis. There were a few articles which identified profitable technical strategies, such as Sweeney and Osler ².

Compared to the dozens of studies of technical analysis in the equity and foreign exchange markets, there are relatively few studies of technical analysis in commodity futures mar-

²Spyros Skouras has compiled an exhaustive bibliography of academic studies through 1998. It is available at <http://www.santafe.edu/spyros/tabiblio.htm>

kets. Lukac and Brorsen, and Lukac, Brorsen and Irwin (1988, 1988a, and 1989) are the only studies of technical analysis in commodity futures markets.

There is little dispute that technical analysis is very common among practitioners. Oberlechner surveys foreign exchange traders on their use of technical analysis, and finds that “Only a very small minority of foreign exchange traders demonstrate an exclusively fundamental or exclusively chartist overall forecasting approach.” This is consistent with the previous survey research performed by Taylor and Allen, Menkhoff, and Lui and Mole. Brorsen and Irwin find similar results for commodity trading advisors.

Using genetic programming, this paper develops optimal *ex ante* trading rules for various commodity markets. Each trading rule is generated using two sequential futures contracts of identical maturity month, and then tested using the next contract of identical maturity for its out-of-sample performance. These tests reveal that these trading rules are quite capable of forecasting periods of high and low returns. The trading strategies are capable of generating profits, but when transactions costs are included, these profits become negligible.

This article has four sections. The first section explains the use of genetic programs in constructing and optimizing technical trading strategies. The second section discusses the evaluation of futures trading strategies and the data used. The third section presents the results of these rules, while the fourth section offers a summary and conclusion.

1 Genetic Programs, Data-Snooping, and Technical Analysis

Genetic programming is the subdiscipline of evolutionary algorithms in which complex algorithms or programs are built from hierarchies of simple operators; they trace their origins to Koza. These programs are optimized according to a evolutionary process whereby an initial population of random rules is generated, they are evaluated according to some ‘fitness’ function, and then ‘evolved’ through random combination to form a new generation of rules.

The use of genetic programs as technical trading strategies dates to Neely, Weller, and Dittmarr, (NWD) and Allen and Karjalainen (AK). These researchers recognized that genetic programming avoids the data-snooping biases inherent in earlier technical research. In most prior technical analysis research, the performance of common trading rules is evaluated using historical data. However, the fact that these rules are common or popular is *prima facie* evidence that they have been profitable in the past. Evaluating these rules using historical data is thus little more than *ex post* model-fitting.

Genetic programs avoid data-snooping because the rules constructed are drawn from the space populated by combinations of simple arithmetic and logical operators. These rules are mechanically optimized using historical data, and then tested using a different set of historical data. Therefore, although the specification of these rules is dependent upon




Figure 1: Simple Trinary Trading Rule

their historical performance, their merit is judged using data not available during their construction.

The trading rules used in NWD and AK were binary, i.e. they could only indicate two states for the investor. These states were variously mapped to trading positions of long/short, long/neutral, or neutral/short. While binary positions may make sense in equities, they are problematic in futures markets because there is no physical asset being held, and short positions are taken as easily as long positions. Therefore, this article proposes the use of trinary trading rules, in which the rule can indicate long, neutral, or short positions.

Figure 1 is an example of a simple trading rule as they are used in this study. **XR** is a *root node* that requires two *subnodes*, which for this rule are the inequality operators $>$ and $<$. The real values 500 and 10000, as well as the data **VOL** (volume) and **CPR** (closing price) are *terminal nodes*, nodes which do not have subnodes. **XR** is a trinary operator whose state is a function of the states of its subnodes (in this case, the subnodes are $>$ and $<$), as displayed in table 1, where long, neutral, and short positions are indicated by 1, 0, and -1, respectively. Rule 1 indicates a long position should be taken if the closing price is above 500 but the volume is greater than 10,000; a short position should be taken if the reverse is true, and no position should be taken if both or neither are true.

Table 1: State of **XR** given the subnode states

| XR | Subnode 1 | Subnode 2 |
|-----------|--------------|--------------|
| 0 | TRUE | TRUE |
| 1 | TRUE | FALSE |
| -1 | FALSE | TRUE |
| 0 | FALSE | FALSE |

The choice of nodes in building the genetic programs is similar to those used in NWD and AK. Terminal nodes (those that take no arguments) may be real $[-10,10]$, boolean (**TRUE**, **FALSE**), or return price data: **OPR**, **HPR**, **LPR**, **CPR**, **VOL**, and **OI** represent the opening, high, low, and closing price, and the daily volume and open interest. Function nodes can be the arithmetic operations, $+$, $-$, \times , and \div , boolean operators, **IF-THEN-ELSE**, **AND**, **OR**, **NOT**, inequalities, $<$, $>$, square, square root, and the 1-norm (distance). Additionally, four functions are included that operate on lagged data, each of which requires two arguments,

Table 2: Nesting of Common Technical Indicators within Functional/Terminal Node Sets.

| Technical Indicator | Nested in Node Sets | |
|-------------------------|---------------------|---------|
| | AK | Current |
| Trend Lines | - | X |
| Support/Resistance | X* | X |
| Channel Line | - | X |
| Percentage Retracements | X | X |
| Speedlines | X | X |
| Gaps | - | X |
| Head and Shoulders | - | - |
| Double Tops/Bottoms | - | - |
| Triangles | - | - |
| Moving Average | X | X |
| Envelopes | X | X |
| Bollinger Bands | X* | X |
| Momentum | X | X |
| RSI | X | X |
| Stochastics | - | X |
| % R | - | X |
| MACD | X | X |
| Candlesticks | - | X |

* Although these indicators can be based only on closing prices, high and low prices are most commonly used.

a data series (OPR, HPR, LPR, or CPR) and a real value, k , which indicates the number of prior observations over which to operate. LAG returns the k -day lagged price, MIN and MAX return the minimum and maximum values over the k days periods, and AVG returns the k -day average. MND and MXD are similar to MIN and MAX except that they return the number of days since the lowest (highest) value in the last k days. Table 2 lists the first eighteen technical trading indicators in Murphy. The set of rules that can be constructed using the operator set in this article encompasses most common technical rules, and is significantly expanded from NWD and AK.

The evolutionary process used to generate optimal trading rules is the defining characteristic of genetic programming. To start the process, a population of rules is randomly generated. Each of these rules is evaluated for its ‘fitness’—such as high profitability or low risk. With a probability proportional to each rule’s fitness rank in the population, rules are chosen to participate in genetic operations, such as recombination, and the resulting rules constitute the next generation of rules. This three-step process (evaluate, select, operate) is repeated until convergence or the maximum number of iterations is reached.

In this study, only two genetic operators are used to create rules. In *reproduction*, rules




Figure 2: Recombination of Two Technical Trading Rules

from the parent generation are inserted into the child generation unchanged. In *recombination*, two parent rules are chosen, and sub-trees are randomly chosen from each parent rule and exchanged. Figure 1 shows the recombination of two parent rules into two child rules. While many other genetic operations have been proposed, reproduction and recombination are the two most common, and additional rules typically offer little benefit. (Koza)

Because rules are selected for operation based upon their fitness, the specification of the fitness measure is crucial for the success of genetic programming. Two fitness measures are used to generate the rules in this study, gross profitability and the ratio of profitability to maximum intermediate loss. These two criteria will be explained in the next section.

Initially, 10,000 randomly-generated rules are created. The 1,000 fittest are retained to make up the first generation. Each successive generation consists of the fittest rule from the previous generation, 99 randomly-chosen rules are inserted unaltered (reproduction), and the remaining 900 are the product of recombination of randomly-chosen pairs. Analogously to the evolutionary process, rules are not truly randomly chosen. Instead, the probability that a rule will be chosen for insertion or recombination is a function of its fitness. Specifically, the probability is a function of a rule's rank within the population,

$$p_i = \frac{r_i^3}{\sum_j r_j^3} \quad (1)$$

where p_i is the probability that the i^{th} rule will be chosen, where i is the ordinal rank of the rule, with $i = N$ the most fit, and $i = 1$ the least fit.

In order to prevent over-fitting, the rules are generated using two sets of futures price data,

as in Allen and Karjalainen. Rules are evaluated for selection and operation based upon their fitness in ‘training’ data, which are one year’s worth of prices for a given commodity futures contract of a given expiration month. After each generation is evaluated using the training data, the fittest rule is applied to the ‘selection’ data, which is also one year’s price data, of the same maturity month as the training data, but from the following maturity year. If this rule is fitter than the previous rules evaluated with selection data, it is retained.

Because GP cannot guarantee convergence, either locally or globally, the quality of a solution is a monotonic function of its computational cost; as larger populations of larger rules are allowed to evolve longer, the probability of convergence increases. Balancing this need is the time required for estimation. The population size is 1000 rules, each of which is constrained to 100 nodes. In the initial rule generation, the rules are constrained to be no more than seven levels deep, but in recombination, the rules can grow to be 16 levels deep. To further improve the results, ten optimizations are performed over each set of training/selection data, differing only in the seed value to the random number generator, and the best rule of ten is used in subsequent testing.

Finally, the rule that emerges from the testing/selection process is applied to out-of-sample data. The out-of-sample evaluation uses one year of prices of the same maturity month as the testing and selection data, but from the following contract year.

2 Trading Strategy Evaluation

Net profits are the simplest and most common measure of the usefulness of a trading strategy. The leveraged nature of futures contracts makes the use of simple return-based measures of performance more difficult, as it is unclear what denominator should be used in computing the return. One could assume that no leverage is possible, although this seems a very strong assumption, especially as leverage is frequently cited as an advantage of futures markets. Alternatively, one could use the margin requirement as the denominator. This is also problematic, as US Treasury Bills can be pledged as collateral, meanwhile still accruing interest for the futures-holder, which reduces the forgone interest of holding futures to zero.

For these reasons, simple profitability is used instead of returns. The fitness measure used is

$$\pi = \sum_t^T (p_{t+1} - p_t) I_t - \phi \text{abs}(I_t - I_{t-1}) \quad (2)$$

where $I_t \in [-1, 0, 1]$ is the trading position at time t and ϕ is the transactions cost. As suggested by Neely, Weller and Dittmar, higher transactions costs discourage rules which over-trade, which may be a symptom of over-fitting. They recommend using a transaction cost that is higher than otherwise may be realistic for training and selection, and a more realistic rate for out-of-sample testing. Therefore, transactions costs of \$25 per round-trip are used for training and selection, and \$6.25/round-trip are used in out-of-sample testing,

approximating the commission level of a large trading firm.

Profits are not the only measure of a successful trading strategy. Other useful criteria might be low variance or small intermediate losses. The second fitness measure considered here uses the concept of drawdown, or intermediate loss as a proxy of variability, to measure a strategy's fitness. The concept of drawdown is especially relevant to futures trading strategies, as the maximum margin requirement is a monotonic function of drawdown. The drawdown of a strategy is defined as the difference between the highest intermediate profit of the strategy and its current value, or, if $\Pi_T = \sum_t \pi_t$, then drawdown is

$$\delta_T = \max_{t=1,\dots,T} \left(\max_{\tau=1,\dots,t} (\Pi_\tau) - \Pi_t \right) \quad (3)$$

The second fitness measure is the ratio of profit to maximum drawdown, π/δ .

While an almost limitless number of fitness functions could be conceived, these two represent computationally-efficient measures that proxy the interests of agents using technical trading strategies.

3 Results

Optimal trading rules are estimated for CBOT Corn, Soybean and Wheat futures. Training is performed using five maturities of data for each year from 1978 through 1998³, yielding 105 rules for each commodity. In order that any seasonal factors may be preserved, selection and testing are performed using prices from subsequent years' contracts, but with the same maturity month, i.e. the first rules generated for each commodity use data from the March 1978 contract for training, the March 1979 contract for selection, and the March 1980 contract for out-of-sample testing. For each contract maturity, the data from the first trading day of the delivery month of the previous calendar year through the final trading day of the full calendar month prior to delivery is used, i.e. for the March 1978 contract, data from 1 March 1977 through 29 February 1978 is used. This method of construction ensures that there is no overlap between the training, selection or testing datasets.

Tables 3, 4, and 5 report the performance of the generated rules for corn, soybeans, and wheat, respectively. The first two sets of results on each table are those of the genetic programming rules and a strategy of purchasing and holding the corresponding contract for the entire period. For all three commodities, the genetic programming method is capable of finding extremely profitable trading rules in-sample. Using the profit-only fitness criteria, π , the daily average in-sample profit of these rules is \$11.13, \$33.77, and \$18.01 for the three commodities. Using the second fitness criteria, π/δ , the profits are smaller, but still greater than the static strategies, and they are achieved with a reduction in average daily volatility ($\bar{\sigma}$) of 61%, 71% and 65% respectively. The rules are also clearly

³For Corn and Wheat, March, May, July, September and December contracts are used. For soybeans, November contracts are used instead of December.

able to discern periods of high returns from those of low and negative returns using in-sample data, as evidenced by comparing the average returns on days on which the rules were long, neutral, or short in tables 3, 4, and 5. These findings are confirmed by the pairwise tests in table 6, in which the profit of the technical rule was statistically greater than that of the static rule for each commodity.

In out-of-sample testing, the performance of the rules is less compelling. Only two of the six commodity/fitness combinations manage to generate positive mean returns, both for wheat. Of these two, only the strategy that uses π/δ as the fitness measure has returns that are statistically greater than zero, and then only at the 10% level. The mean profit of the technical strategies is greater than the static rules in five of the six combinations, but the differences are statistically insignificant for corn and one may question whether the finding for soybeans is the result of the sharp price declines experienced by soybeans during the period of study (notice that the average daily mean return of soybeans during the sample period is -5.7522). Comparing the returns on days of long and short positions reveals that the rules are able to consistently identify periods of high, low, and negative returns out-of-sample for only the case of π/δ applied to wheat futures, for which the rules are able to distinguish between these periods at 5% or greater levels of significance.

According to table 6, the rules generated in the wheat market do produce profits that are statistically larger than a static long position, but only the π/δ rules are statistically different from 0, and then only barely. Both fitness measures are significantly higher than the static long position. Profitable exploitation of the difference between the technical strategy and the static long position requires maintaining a static short position in addition to the position indicated by the technical strategy. This combined strategy increases transactions costs only minimally (one round-turn per year), and would not change the statistical significance of the wheat rules, or much alter their indicated profitability, of \$1250/contract, with an average historical drawdown of \$1600.

Using the profit/drawdown measure, the wheat rules also appear to be able to differentiate periods of high returns from low returns, and from high returns to negative returns at the 5% level.

For the corn and soybean markets, these results confirm prior findings in the equity and foreign exchange markets that technical trading rules do not appear to be able to generate economic profits in the presence of transactions costs. Rules generated with genetic programming are clearly able to discern between periods of high and low profits in-sample, but fail to do so in out-of-sample testing. In both in-sample and out-of-sample applications, these rules are capable of reducing the volatility of returns, but it remains unclear whether they are able to achieve this more successfully than existing methods. The evidence in the wheat market is less clear. The rules generated for the wheat market using the profit/drawdown fitness measure are able to discern between periods of high and low returns and very high levels of significance, and are able to generate economic profits, though the statistical evidence for profitability is weaker.

4 Summary and Conclusion

The prevalence of technical analysis in commodity markets is a mystery. As a method of generating profits, it directly contradicts weak-form efficiency. While many explanations of technical analysis have been offered, none have provided any reason to expect sustained economic profitability of technical methods.

Much of the prior research into technical analysis has been hampered by data-snooping biases, introduced when popular technical methods are evaluated using historical data. In order to avoid data-snooping, this paper uses a genetic programming algorithm to generate optimal technical trading rules for three agricultural futures markets, which are then tested using out-of-sample data.

While these rules are quite successful at identifying periods of high, low, and negative returns *ex post*, the rules for corn and soybeans are not capable of generating profits in the presence of transactions costs when applied to out-of-sample data. The rules do produce higher mean returns when compared to a static long strategy, but the difference is not statistically significant. The rules generated from wheat futures are capable of generating small but significant profits when compared to a static long position. Rules generated to maximize the profit/drawdown ratio are capable of reducing the daily variance of profits compared to a static long portfolio, but it is unclear whether the technical methods used here are any more or less useful than conventional statistical methods in prediction of volatility.

The results of this study can, at best, only be viewed as a lower bound to the profitability of technical analysis. The function set used in this study does not encompass all technical indicators. Because genetic programming is a stochastic search method, the rules used in evaluation in this study are not guaranteed to be globally, or even locally, optimal. Therefore, rules that lie within the domain of this study may exist that are superior to this study's 'optimal' rules but were not identified by the optimization process. Further, the rules used did not even incorporate basic investment management practices, such as stop-loss orders. Finally, the fitness functions used in this study do not incorporate a risk-return tradeoff in describing the desirability of a given trading strategy. Each of these factors could contribute to the lack of support for the use of technical analysis indicated by this study.

However, the evaluation method used in this study also made the relatively strong assumptions that period t closing prices can be used in the period t trading decision, and that trading takes place at the period t closing price.

While each of the above assumptions provide avenues for future research, the two most interesting are the specification of some form of a mean-variance utility function as the fitness criteria, as well as a comparison of the ability of technical analysis to forecast volatility, possibly in combination with returns. Finally, the results of the wheat futures should be explored more carefully; while the profitability of the rules generated is not statistically different from 0, it is significantly different from a static long futures position, and these rules do appear to be capable of generated a small profit.

Table 3: Summary of Technical Trading Rules Applied to CBOT Corn Futures

| | Selection Data | | Out of Sample Data | |
|---------------------------|----------------|--------------|--------------------|--------------|
| | π | π/δ | π | π/δ |
| Dynamic Trading Strategy | | | | |
| μ | 11.1298 | 7.0476 | -2.6107 | -0.9554 |
| σ_μ | 9.6512 | 4.7448 | 8.6063 | 5.4917 |
| $\bar{\sigma}$ | 147.8555 | 81.5111 | 143.1905 | 82.2999 |
| Static Long Position | | | | |
| μ | -2.0090 | | -2.5142 | |
| σ_μ | 12.2683 | | 12.5683 | |
| $\bar{\sigma}$ | 161.1638 | | 161.2642 | |
| Dynamic Strategy, Long | | | | |
| μ | 13.8808 | 33.2119 | -6.4772 | -6.3462 |
| σ_μ | 23.3850 | 78.9139 | 21.8542 | 45.3998 |
| $\bar{\sigma}$ | 184.4950 | 192.6476 | 162.7179 | 153.1524 |
| % Days | 0.3386 | 0.1027 | 0.3311 | 0.0964 |
| Dynamic Strategy, Neutral | | | | |
| μ | 2.0308 | -1.7468 | -0.9682 | -2.3335 |
| σ_μ | 63.1136 | 15.6761 | 32.2702 | 13.3987 |
| $\bar{\sigma}$ | 152.4209 | 161.7651 | 165.0259 | 160.4383 |
| % Days | 0.1789 | 0.7398 | 0.2043 | 0.7483 |
| Dynamic Strategy, Short | | | | |
| μ | -14.5431 | -25.8632 | -0.2507 | -0.6490 |
| σ_μ | 37.1981 | 38.5606 | 19.5875 | 49.2405 |
| $\bar{\sigma}$ | 147.1356 | 137.3965 | 159.6053 | 172.8634 |
| % Days | 0.4825 | 0.1575 | 0.4647 | 0.1553 |
| All Rules | | | | |
| $n > 0$ | 100 | 104 | 42 | 46 |
| δ | 34.1548 | 11.8690 | 55.9500 | 25.8238 |
| Mean R/T | 23.6095 | 17.5143 | 23.4476 | 17.8000 |
| Max(π) | 44.2729 | 25.5727 | 15.7750 | 9.4124 |
| Min(π) | -24.7768 | -2.4457 | -29.9500 | -33.8000 |

105 rules were generated. π is the profit net of transactions costs, δ is the maximum drawdown (see equation 3), μ is the mean daily return, σ_μ is the standard error of μ , and $\bar{\sigma}$ is the mean of the daily standard deviation of profits across all contracts; μ , σ_μ , and $\bar{\sigma}$ are expressed in \$/day. Transactions costs were applied at \$6.25 per round-trip. Statistics for Long, Neutral and Short positions are averaged across the entire data set.

Table 4: Summary of Technical Trading Rules Applied to CBOT Soybean Futures

| | Selection Data | | Out of Sample Data | |
|---------------------------|----------------|--------------|--------------------|--------------|
| | π | π/δ | π | π/δ |
| Dynamic Trading Strategy | | | | |
| μ | 33.7710 | 11.0623 | -0.2043 | -0.6570 |
| σ_μ | 20.1602 | 7.2439 | 21.1130 | 10.8098 |
| $\bar{\sigma}$ | 398.1078 | 116.8453 | 400.0987 | 161.0494 |
| Static Long Position | | | | |
| μ | -4.7514 | | -5.7522 | |
| σ_μ | 25.2896 | | 25.2956 | |
| $\bar{\sigma}$ | 436.2584 | | 430.8642 | |
| Dynamic Strategy, Long | | | | |
| μ | 37.2273 | 176.4959 | -6.5270 | -24.9057 |
| σ_μ | 53.6884 | 208.5749 | 56.7529 | 154.4124 |
| $\bar{\sigma}$ | 459.6981 | 484.6674 | 436.9672 | 423.9245 |
| % Days | 0.3687 | 0.0187 | 0.3625 | 0.0202 |
| Dynamic Strategy, Neutral | | | | |
| μ | 17.0027 | 0.0264 | -2.2106 | -5.6466 |
| σ_μ | 96.4147 | 29.5923 | 107.8516 | 27.1456 |
| $\bar{\sigma}$ | 469.7569 | 439.5420 | 430.0410 | 422.1686 |
| % Days | 0.1463 | 0.9148 | 0.1404 | 0.8966 |
| Dynamic Strategy, Short | | | | |
| μ | -43.0809 | -120.2938 | -6.0136 | -1.2071 |
| σ_μ | 77.5271 | 228.8401 | 63.5424 | 187.0107 |
| $\bar{\sigma}$ | 409.7164 | 386.3241 | 429.5237 | 526.6813 |
| % Days | 0.4849 | 0.0665 | 0.4972 | 0.0832 |
| All Rules | | | | |
| $n > 0$ | 101 | 105 | 50 | 53 |
| δ | 82.1071 | 10.3095 | 136.4334 | 42.6619 |
| Mean R/T | 34.7429 | 9.5619 | 33.2667 | 10.2667 |
| Max(π) | 76.8105 | 34.9851 | 47.4250 | 22.3855 |
| Min(π) | -45.0697 | 0.3735 | -62.8224 | -52.3810 |

105 rules were generated. π is the profit net of transactions costs, δ is the maximum drawdown (see equation 3), μ is the mean daily return, σ_μ is the standard error of μ , and $\bar{\sigma}$ is the mean of the daily standard deviation of profits across all contracts; μ , σ_μ , and $\bar{\sigma}$ are expressed in \$/day. Transactions costs were applied at \$6.25 per round-trip. Statistics for Long, Neutral and Short positions are averaged across the entire data set.

Table 5: Summary of Technical Trading Rules Applied to CBOT Wheat Futures

| | Selection Data | | Out of Sample Data | |
|---------------------------|----------------|--------------|--------------------|--------------|
| | π | π/δ | π | π/δ |
| Dynamic Trading Strategy | | | | |
| μ | 18.0116 | 8.4777 | 1.0794 | 0.9259 |
| σ_μ | 12.1431 | 5.6898 | 13.2608 | 7.4077 |
| $\bar{\sigma}$ | 211.3358 | 94.4503 | 209.8399 | 112.6147 |
| Static Long Position | | | | |
| μ | -1.7468 | | -3.2982 | |
| σ_μ | 16.4167 | | 15.6400 | |
| $\bar{\sigma}$ | 234.9496 | | 231.3083 | |
| Dynamic Strategy, Long | | | | |
| μ | 23.4439 | 57.4797 | -2.0045 | 7.5704 |
| σ_μ | 33.2704 | 60.4151 | 44.1327 | 46.7534 |
| $\bar{\sigma}$ | 253.3747 | 256.8510 | 246.0227 | 252.7633 |
| % Days | 0.3831 | 0.0796 | 0.3935 | 0.1054 |
| Dynamic Strategy, Neutral | | | | |
| μ | -5.2467 | -2.4332 | 0.6420 | -4.4644 |
| σ_μ | 55.1146 | 15.6063 | 43.5340 | 18.9775 |
| $\bar{\sigma}$ | 244.2229 | 237.9467 | 249.4612 | 228.3863 |
| % Days | 0.1779 | 0.8163 | 0.1539 | 0.7821 |
| Dynamic Strategy, Short | | | | |
| μ | -22.2534 | -41.3724 | -5.7060 | -5.1438 |
| σ_μ | 32.8844 | 79.7937 | 33.1215 | 37.8559 |
| $\bar{\sigma}$ | 214.4732 | 191.9620 | 211.6171 | 233.3540 |
| % Days | 0.4390 | 0.1042 | 0.4526 | 0.1125 |
| All Rules | | | | |
| $n > 0$ | 98 | 103 | 55 | 56 |
| δ | 45.2857 | 12.5262 | 74.5190 | 32.4229 |
| Mean R/T | 30.3714 | 16.8476 | 28.0667 | 17.8952 |
| Max(π) | 44.1968 | 28.2700 | 36.2301 | 24.4250 |
| Min(π) | -25.4216 | -7.7133 | -31.1250 | -19.2978 |

105 rules were generated. π is the profit net of transactions costs, δ is the maximum drawdown (see equation 3), μ is the mean daily return, σ_μ is the standard error of μ , and $\bar{\sigma}$ is the mean of the daily standard deviation of profits across all contracts; μ , σ_μ , and $\bar{\sigma}$ are expressed in \$/day. Transactions costs were applied at \$6.25 per round-trip. Statistics for Long, Neutral and Short positions are averaged across the entire data set.

Table 6: Summary of Pairwise Tests

| | Selection Data | | Out of Sample Data | |
|------------------------------|----------------|--------------|--------------------|--------------|
| | π | π/δ | π | π/δ |
| Corn | | | | |
| $\pi > 0$ | 11.8169*** | 15.2200*** | -3.1084 | -1.7827 |
| $\pi > \pi^{static}$ | 8.6250*** | 7.0552*** | -0.0649 | 1.1645 |
| $\pi^{long} > \pi^{neutral}$ | 1.8041** | 4.4524*** | -1.4484 | -0.8687 |
| $\pi^{long} > \pi^{short}$ | 6.6288*** | 6.8921*** | -2.1740 | -0.8716 |
| Soybeans | | | | |
| $\pi > 0$ | 17.1650*** | 15.6483*** | -0.0992 | -0.6228 |
| $\pi > \pi^{static}$ | 12.2051*** | 6.1597*** | 1.7254** | 1.8980** |
| $\pi^{long} > \pi^{neutral}$ | 1.8779** | 8.5837*** | -0.3629 | -1.2587 |
| $\pi^{long} > \pi^{short}$ | 8.7263*** | 9.8220*** | -0.0618 | -1.0013 |
| Wheat | | | | |
| $\pi > 0$ | 15.1991*** | 15.2676*** | 0.8341 | 1.2808* |
| $\pi > \pi^{static}$ | 9.9151*** | 6.0300*** | 2.1876** | 2.5012*** |
| $\pi^{long} > \pi^{neutral}$ | 4.5666*** | 9.8388*** | -0.4375 | 2.4440*** |
| $\pi^{long} > \pi^{short}$ | 10.0099*** | 10.1207*** | 0.6874 | 2.1657** |

* indicates significance at the 10% level, ** indicates significance at the 5% level, and *** indicates significance at the 1% level. Test statistics used are

$$z = \frac{r_1 - r_2}{\sqrt{\frac{\sigma_1^2 + \sigma_2^2}{T}}}$$

where $z \sim N(0, 1)$.




Figure 3: Profitability of π rules and static long positions, CBOT Corn Futures




Figure 4: Profitability of π/δ rules and static long positions, CBOT Corn Futures




Figure 5: Profitability of π rules and static long positions, CBOT Soybean Futures




Figure 6: Profitability of π/δ rules and static long positions, CBOT Soybean Futures




Figure 7: Profitability of π rules and static long positions, CBOT Wheat Futures




Figure 8: Profitability of π/δ rules and static long positions, CBOT Wheat Futures

References

- Alexander, S. S.**, “Price Movements in Speculative Markets: Trends or Random Walks, No. 2,” in Paul Cootner, ed., *The Random Character of Stock Market Changes*, MIT Press, 1964, pp. 338–372.
- Allen, Franklin and Risto Karjalainen**, “Using genetic algorithms to find technical trading rules,” *Journal of Financial Economics*, 1999, *51*, 245–271.
- Brorsen, B. Wade and Scott H. Irwin**, “Futures Funds and Price Volatility,” *Review of Futures Markets*, 1987, *7*, 118–135.
- Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay**, *The Econometrics of Financial Markets*, Princeton, NJ: Princeton University Press, 1997.
- Fama, Eugene and Marshall Blume**, “Filter Rules and Stock Market Trading,” *Journal of Business*, 1966, *3*, 226–241.
- Koza, John R.**, *Genetic Programming: On the Programming of Computers by Means of Natural Selection*, Cambridge: MIT Press, 1992.
- Lui, Yu Hon and David Mole**, “The Use of Fundamental and Technical Analyses by foreign exchange dealers: Hong Kong evidence,” *Journal of International Money and Finance*, 1998, *17*, 535–545.
- Lukac, Louis P. and B. Wade Brorsen**, “A Comprehensive Test of Futures Market Disequilibrium,” *The Financial Review*, 1990, *25* (4), 593–622.
- , — , and **Scott H. Irwin**, “Similarity of Computer Guided Technical Trading Systems,” *Journal of Futures Markets*, 1988, *8*, 1–13.
- , — , and — , “A Test of Futures Market Disequilibrium using Twelve Different Technical Trading Systems,” *Applied Economics*, 1988, *20*, 623–639.
- , — , and — , “The Usefulness of Historical Data in Selecting Parameters for Technical Trading Systems,” *Journal of Futures Markets*, 1989, *9*, 55–65.
- Malkiel, Burton G.**, *A Random Walk Down Wall Street*, 4th ed., New York: Norton, 1985.
- , “Efficient Markets Hypothesis,” in P. Newman, M. Milgate, and J. Eatwell, eds., *New Palgrave Dictionary of Money and Finance*, London: Macmillan, 1992.
- Menkhoff, L.**, “Examining the use of technical currency analysis,” *International Journal of Finance and Economics*, 1997, *2*, 307–318.
- Murphy, John J.**, *Technical Analysis of the Financial Markets*, revised ed., New York: New York Institute of Finance, 1999.

- Neely, Christopher J., Paul Weller, and Rob Dittmar**, “Is technical analysis in the foreign exchange market profitable? A genetic programming approach,” *Journal of Financial and Quantitative Analysis*, 1997, 32 (4), 405–426.
- Oberlechner, Thomas**, “Importance of Technical and Fundamental Analysis in the European Foreign Exchange Market,” *International Journal of Finance and Economics*, 2001, 6, 81–93.
- Osler, Carol**, “Identifying Noise Traders: The Head-and-Shoulders Pattern in U.S. Equities,” 1998. Staff Report No. 42, Federal Reserve Bank of New York.
- Sweeney, Richard J.**, “Some New Filter Rule Tests: Methods and Results,” *Journal of Financial and Quantitative Analysis*, 1988, 23 (3), 285–300.
- Taylor, Mark P. and Helen Allen**, “The Use of Technical Analysis in the Foreign Exchange Market,” *Journal of International Money and Finance*, 1992, 11 (3), 304–314.