

Performance Analysis of a Counter-intuitive Automated Stock-Trading Agent

Ronggang Yu and Peter Stone

Department of Computer Sciences

The University of Texas at Austin

<http://www.cs.utexas.edu/~{ryu,pstone}>

ABSTRACT

Autonomous trading in stock markets is an area of great interest in both academic and commercial circles. A lot of trading strategies have been proposed and practiced from the perspectives of Artificial Intelligence, market making, external data indication, technical analysis, etc. This paper examines some properties of a counter-intuitive automated stock-trading strategy in the context of the Penn-Lehman Automated Trading (PLAT) simulator [1], which is a real-time, real-data market simulator. While it might seem natural to buy when the market is on the rise and sell when it is on the decline, our strategy does exactly the opposite. As a result, we call it the *reverse* strategy. The reverse strategy was the winner strategy in the first and second PLAT live competitions. In this paper, we analyze the performance of the reverse strategy. Also, we suggest ways to control the risk of using the reverse strategy in certain kinds of markets.

1. INTRODUCTION

Many major stock markets are electronic. The NASDAQ is a distributed trading system completely run through networked computers. It allows customers' best bids and offers to be displayed and represented on the NASDAQ by their brokers or through ECNs (Electronic Communication Networks). ECNs such as Archipelago [2], Bloomberg [3] and Island [4] allow customers to display their orders and also allow customer orders to be traded with each other.

The ECNs are easy to use, and available to everyone, yet they are not platforms for customers to test their trading strategies. There are a lot of simulators developed for users to use without risking money in real markets. The Stock Market Game [5] is a simulator that enables participants to discover the risks and rewards involved in decision-making. Virtual Stock Exchange [6] is another simulator that participants can use to build and manage a virtual portfolio. The Penn-Lehman Automated Trading (PLAT) simulator

uses real-world, real-time stock market data available over modern ECNs. It is the first simulator to incorporate complete order book information, thereby allowing it to "match" agent orders with real-world orders and simulate the resulting effects on the market. It also provides a lot of APIs so that the participants can program their strategies and trade with other agents and outside markets automatically.

Using these and other simulators, as well as experiments in the real world, many researchers have studied trading strategies from the perspectives of Artificial Intelligence [7, 8, 11], neural networks [9], technical analysis [10], etc.. However, experimenting in the real world is expensive and most simulators differ significantly from real markets such that strategies successful in simulation may not be appropriate for real markets. PLAT is among the most realistic simulators because it both includes real-time, real-world data, and realistically models the effects of the agents' own trades on the market. Thus, we think some interesting, potentially applicable conclusions can be reached based on agent experiments in PLAT.

In this paper, we present a successful, counter-intuitive automated stock-trading strategy that we have implemented and tested in PLAT. While it might seem natural to buy when the market is on the rise and sell when it is on the decline, our strategy does exactly the opposite. As a result, we call it the *reverse* strategy.

The remainder of this paper is organized as follows. In section 2, we give a brief introduction to the PLAT simulator. In section 3, we describe two opposite trading strategies, including our reverse strategy. A theoretical analysis of the reverse strategy is provided in detail in section 4. In section 5, we present detailed empirical results, including results from the the first and second PLAT live competitions in which the reverse strategy was the winner strategy. Finally, suggestions for improving the performance of the reverse strategy and further work are discussed in section 6.

2. THE PLAT SIMULATOR

The PLAT simulator uses real-world, real-time stock market data for automated trading. It frequently queries the Island ECNs web-site to get the most recent stock prices and buy and sell order books. The simulator then matches the buy orders and sell orders. The orders can be from Island or from trading agents. The simulator also computes the prof-

its and losses of each connected trading agent in real time.

PLAT is equipped for testing strategies on historical data and also for running live competitions. The live competition starts and ends at the same time with normal trading sessions of the NASDAQ. The simulator supports limit orders only. A limit order is a request to buy or sell shares of a stock at a specified price. In the simulation, the best ask price is the lowest price any seller (either trading agents or outside market customers) has declared that they are willing to accept; the best bid price is the highest price any buyer has declared that they are willing to pay. If a new buy order has bid price equal to the best ask price or a new sell order has ask price equal to the best bid price, the order will be matched in the amount of the maximum available shares and the trade is executed. If a bid price is higher than the ask price, the trading price is the average of these two prices. If orders cannot be matched immediately, they are kept in the queue to wait for possible future matches.

Currently, The PLAT simulator is hardwired to Microsoft Stock (Symbol: MSFT). Trading agents in the simulation can buy or sell MSFT stocks with limit orders. There are two types of sales supported in the simulation: long and short. Long sales are what we normally think of when we think of selling, that is, sales of stocks owned by the seller. Short sales are sales of stocks borrowed by the agents from the simulator. Trading agents can also borrow money from the simulator without any interest, however, the same amount of money will be deducted from the agent's simulated cash account. The cash in the account can be negative or positive. The value that a trading agent has in the simulation is calculated in real time by the formula: $value = cash + holdings * current Price$. The cash and holdings are set to 0 at the beginning of the simulation.

The PLAT simulator differs from real markets in several ways:

- There's some time lag in the simulator either due to the server being overloaded or the lag of the data available from the Island ECNs. Resourceful trading agents could potentially derive some advantage from gaining access to a faster real-time data source, but to our knowledge nobody has done so.
- There are no commission or tax charges in the PLAT simulation. Transactions can be executed without any fees. In the real market, too many transactions will increase the overhead cost.
- The trading in PLAT is fully automated, meaning that once strategies are fixed, participants cannot intervene during the trading day.

Our proposed strategies utilize the fact that there are no commission or tax charges in the simulation, and places orders as frequently as possible. In the following sections, we describe two agent strategies and analyze under what market conditions they can make profits.

3. AGENT STRATEGIES

Our initial *basic* strategy is as follows. At any time during the simulation, if the stock price goes up, it places a buy order; and if the stock price goes down, it places a sell order.

The motivation is that a price rise indicates likely further price rises. However, initial testing of the basic strategy revealed that it lost money more often than it gained. As a result, we decided to flip the buy and sell order conditions. We call the resulting strategy the *reverse* strategy since it does exactly the opposite of the basic strategy. Table 1 and Table 2 show the pseudo-codes for the basic and reverse strategies:

```

while time permits:
lastPrice ← getLastPrice();
currentPrice ← getCurrentPrice();
if currentPrice > lastPrice
    placeOrder(BUY, currentPrice, volume);
elseif currentPrice < lastPrice
    placeOrder(SELL, currentPrice, volume);

```

Table 1: The Basic Strategy

```

while time permits:
lastPrice ← getLastPrice();
currentPrice ← getCurrentPrice();
if currentPrice > lastPrice
    placeOrder(SELL, currentPrice, volume);
elseif currentPrice < lastPrice
    placeOrder(BUY, currentPrice, volume);

```

Table 2: The Reverse Strategy

The reason that the basic strategy will lose money and the reverse strategy will make profits in most kinds of market is that stock market prices are not constant and in fact undergo frequent changes in direction, rather than moving consistently in one direction. The basic strategy will make profits if the changes are consistent. However, on most trading days, there are a lot of small spikes in the stock price chart. Figure 1 shows such an example. Under these conditions, we expect to make profits with the reverse strategy, as described in the next section.



Figure 1: MSFT Stock Price: Dec. 20, 2002

4. ANALYSIS OF STRATEGIES

For any fixed strategy (other than the degenerate do-nothing strategy), there are some market conditions in which it will make money and some under which it will lose money. In this section, we analyze the expected profits of our proposed strategies based on what we think are realistic, though simplified, assumptions about price trajectories. We start with a very restricted model, but then gradually relax our assumptions to arrive at a fairly general, predictive model.

In our initial analysis, we assume that the price of MSFT oscillates consistently in a wave pattern around a constant price p . We call the microvariation of the price (the wave's amplitude) during the trading day δ . We also assume that all orders placed in buy queues and sell queues will be matched at some point during the trading day (perhaps immediately). In our analysis, we assume that the start price is $p - \delta$. The start price is not crucial to the analysis since if it is not $p - \delta$, we can ignore the first several ticks until we have the price $p - \delta$. Figure 2(a) illustrates such a price trend. Tables 3 and Table 4 show the calculations of the changes in holdings, cash, and value for the basic and reverse strategies in one cycle. Each time we place a sell order or buy order, the volume size is v .

time	price	action	holding	cash	value
t_0	$p - \delta$	—	0	0	0
t_1	p	buy	v	$-vp$	0
t_2	$p + \delta$	buy	$2v$	$-2vp - v\delta$	$v\delta$
t_3	p	sell	v	$-vp - v\delta$	$-v\delta$
t_4	$p - \delta$	sell	0	$-2v\delta$	$-2v\delta$

Table 3: The basic strategy over one cycle.

time	price	action	holding	cash	value
t_0	$p - \delta$	—	0	0	0
t_1	p	sell	$-v$	vp	0
t_2	$p + \delta$	sell	$-2v$	$2vp + v\delta$	$-v\delta$
t_3	p	buy	$-v$	$vp + v\delta$	$v\delta$
t_4	$p - \delta$	buy	0	$2v\delta$	$2v\delta$

Table 4: The reverse strategy over one cycle.

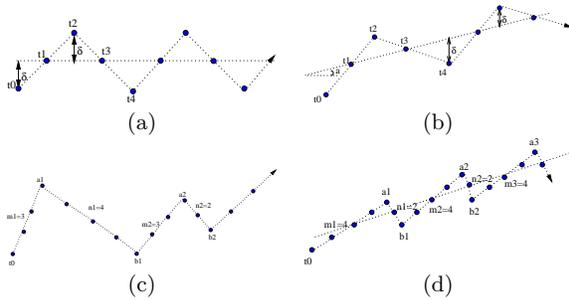


Figure 2: Different price trend models: (a) shows the initial model in which the price oscillates consistently in a wave pattern around a constant price. (b) extends the model to allow for the price to oscillate around a line with non-zero slope. (c) extends the model further to show an extended cycle frequency. (d) shows a special case of (c) that the price still oscillates around a line.

From Table 3 and Table 4, we can see that the basic strategy loses $\$2v\delta$ in one cycle. After k cycles, it loses $\$2vk\delta$. On the other hand, the reverse strategy will earn $\$2vk\delta$: the value increases linearly in time.

Although the preceding calculation relies on a very restrictive model of market dynamics, it is a pattern that seems

to be reasonably representative of reality. In the remainder of this section, we consider several relaxations to our initial assumptions. In particular

1. the price may oscillate around a line with non-zero slope;
2. δ may differ from cycle to cycle;
3. the price may go up repeatedly then down repeatedly rather than staying within δ of p ;

We investigate the performance of the reverse strategy under all of these circumstances.

4.1 Relaxation 1: the price oscillates around a line with non-zero slope (linear model)

As is apparent from Figure 1, it is much more common that the price oscillates around a line with a non-zero slope than that it oscillates around a constant price. We define the price trend as a line: $y = ax + b$. a is the slope, while x is the time elapsed from the start of the simulation, which is measured by the number of simulation ticks. If the tick number is odd, price = y . If it is divisible by 4, price = $y - \delta$. If it is even, but not divisible by 4, price = $y + \delta$. Figure 2(b) illustrates such a price trend. Table 5 shows the calculation of the reverse strategy. In the simulation, δ is always positive, while a can be positive or negative. Since the basic strategy does the exact opposite of the reverse strategy, we omit the calculation.

time	price	action	holding	cash	value
t_0	$b - \delta$	—	0	0	0
t_1	$a + b$	sell	$-v$	$v(a + b)$	0
t_2	$2a + b + \delta$	sell	$-2v$	$v(3a + 2b + \delta)$	$v(-\delta + a)$
t_3	$3a + b$	buy	$-v$	$v(b + \delta)$	$v(\delta - 3a)$
t_4	$4a + b - \delta$	buy	0	$2v(\delta - 2a)$	$2v(\delta - 2a)$

Table 5: The Calculation of the Reverse Strategy: Linear Model

There is one additional assumption in Table 5's calculation: $\delta > |a|$, guaranteeing that the price actually oscillates around the line, rather than moving consistently in one direction. After k cycles, the value of the reverse strategy is $\$2vk(\delta - 2a)$. Combining this formula with the preceding assumption, we see that if $-\delta < a < \delta/2$, the reverse strategy makes a profit in the linear model; otherwise it loses money. Note that if $a = 0$, the result degenerates to that of the constant price model, as expected. Note also that this analysis can be easily extended to the case where the underlying price trajectory is *piece-wise linear*: each line segment can be considered independently.

4.2 Relaxation 2: microvariation changes from cycle to cycle

To this point, we have assumed that δ remains constant throughout the analysis. Here, we relax this assumption by assuming only that δ remains constant throughout an individual cycle, but that it can change from cycle to cycle.

Notice that in the linear model, after each cycle, the holdings is 0. Thus the effects of each cycle are independent. Numbering the cycles sequentially, let δ_i be the microvariation in cycle i and $\bar{\delta}$ be the average of all the δ_i s over the

course of the trading day. Then after k cycles, the value is $\sum_{i=1}^k 2v(\delta_i - 2a) = 2vk(\bar{\delta} - 2a)$. Thus, if $-\bar{\delta} < a < \bar{\delta}/2$, the reverse strategy makes a profit in this model.

4.3 Relaxation 3: Extended cycle frequency

In relaxations 1 and 2, we assumed that the price returns to the underlying linear trajectory every two simulation ticks and that a cycle always lasts exactly 4 ticks. That is, the price never moves in the same direction for more than 2 ticks in sequence. However, as is apparent from the last part of Figure 1, the price cycle can extend longer than 4 ticks in practice. Here we analyze the case in which the price moves in the same direction for several ticks in a row, but still always by the same δ . Figure 2(c) illustrates such a price trend.

Under this model, the price trend in a trading day can be split into a lot of up trends and down trends. We analyze the performance of the reverse strategy in up trends and down trends separately, then combine them. We assume initial holdings are x for the up trend, and y for the down trend. Table 6 and Table 7 show the holdings and change in value of the reverse strategy in one up trend with m ticks and one down trend with n ticks. Note the last column of these two tables are the change of the value from the last action, not the absolute value.

time	current Price	action	holding	Δ value
t_0	p	—	x	—
t_1	$p + \delta$	sell	$x - v$	$x\delta$
t_2	$p + 2\delta$	sell	$x - 2v$	$(x - v)\delta$
t_3	$p + 3\delta$	sell	$x - 3v$	$(x - 2v)\delta$
...
t_m	$p + m\delta$	sell	$x - mv$	$(x - (m - 1)v)\delta$

Table 6: The reverse strategy during an up trend

From Table 6, if the stock price increases m times, $\Delta\text{holding} = -mv$, and $\Delta\text{value} = \sum_{i=0}^{m-1} (x - iv)\delta = mx\delta - \frac{m(m-1)}{2}v\delta$.

time	current Price	action	holding	Δ value
t_0	p	—	y	—
t_1	$p - \delta$	buy	$y + v$	$-y\delta$
t_2	$p - 2\delta$	buy	$y + 2v$	$-(y + v)\delta$
t_3	$p - 3\delta$	buy	$y + 3v$	$-(y + 2v)\delta$
...
t_n	$p - n\delta$	buy	$y + nv$	$-(y + (n - 1)v)\delta$

Table 7: The reverse strategy during a down trend

From Table 7, if the stock price decreases n times, $\Delta\text{holding} = nv$, and $\Delta\text{value} = \sum_{i=0}^{n-1} -(y + iv)\delta = -ny\delta - \frac{n(n-1)}{2}v\delta$. Table 6 and Table 7 show that if the price is continuously decreasing or increasing, our reverse strategy loses money quadratically as time goes on. However, in the extended frequency model, there is still a wave-like price cycle. That is, the up trends and down trends are connected together. Table 8 shows the reverse strategy's value assuming that the price increases m_1 times until tick number a_1 , then decreases n_1 times until tick number b_1 , and so on.

If the price goes up and down in this manner for k cycles, the final value is:

time	holding	Δ value
t_0	0	0
t_{a_1}	$-m_1v$	$-\frac{m_1(m_1-1)}{2}v\delta$
t_{b_1}	$-m_1v + n_1v$	$-m_1(-m_1)v\delta - \frac{n_1(n_1-1)}{2}v\delta$
t_{a_2}	$-m_1v + n_1v - m_2v$	$m_2(-m_1 + n_1)v\delta - \frac{m_2(m_2-1)}{2}v\delta$
t_{b_2}	$-(m_1 + m_2)v + (n_1 + n_2)v$	$-n_2(-m_1 - m_2 + n_1)v\delta - \frac{n_2(n_2-1)}{2}v\delta$
...

Table 8: The Calculation of the Reverse Strategy: Extend Cycle Frequency Model

$$\begin{aligned} \text{value} = & v\delta \left[\sum_{i=2}^k m_i \left(\sum_{j=1}^{i-1} (n_j - m_j) \right) - \sum_{i=1}^k n_i \left(\sum_{j=1}^{i-1} (n_j - m_j) \right) \right. \\ & \left. - m_i \right] + v\delta \left[- \sum_{i=1}^k \frac{(m_i^2 - m_i)}{2} - \sum_{i=1}^k \frac{(n_i^2 - n_i)}{2} \right] \end{aligned}$$

In a market where the price trend is biased, meaning that the number of price increases is larger than the number of decreases or vice versa, the reverse strategy will lose money. To see that, let $n_i = m_i + u$, $n_i = n_j$, and $m_i = m_j = m$, for all i, j . Then, the above expression simplifies to:

$$\text{value} = vk\delta \left[m + \frac{u}{2} - \frac{1}{2}u^2k \right]$$

Thus, whenever u is non-zero the reverse strategy loses money quadratically in time. Figure 2(d) illustrates such a price trend. However, if $u = 0$, i.e., the price trend is not biased, then the reverse strategy earns $vk\delta$. Note that if $m = 2$ and $u = 0$, the value is exactly the same as that in the constant price model, as expected.

4.4 Arbitrary price trajectories

The above analyses have all made two general assumptions. First, they have assumed specific, parameterized price trajectories. Second, they have assumed that all orders are matched at some point during the trading session. In our experiments, we can isolate the impact of these two assumptions by considering a model that only makes the second assumption. That is, we can run the reverse strategy against a stored, complete price trajectory under the assumption that all orders are fully matched. The difference between the calculated value and the value in the corresponding real-time run indicates the extent to which the full-matching assumption influences our predictions.

5. PLAT LIVE COMPETITIONS AND EXPERIMENTAL RESULTS

In this section, we present detailed empirical results demonstrating the behavior of the reverse strategy under real market conditions. First, we describe its performance in the first and second PLAT live competitions that included several agents in addition to our own. Then we present several experiments run on historical data with only our own agent participating.

The first PLAT live competition was held on Dec. 18, 2002. The participants were 23 students with strong interests in the project. 11 teams were formed and divided into two pools: red pool and blue pool. We were one of six teams in the blue pool. In order to limit the divergence of the

simulation from the real market, agents were restricted to maintain holdings within 100,000 of 0 (positive or negative). In the competition, we set our trade volume to 1000 shares per order. This volume size was chosen without detailed experimentation on the one hand to be as large as possible, but on the other hand in the hopes that it was small enough the orders would be matched fully. We also introduced a holding control mechanism to avoid violating the 100,000 share limit. When the holdings exceed 85,000 or are lower than -85,000, the reverse strategy stops placing buy orders or sell orders respectively until the holdings are back to normal level. If the absolute value of holdings exceeds 95,000, the reverse strategy cancels all the orders placed in the buy queue or the sell queue. Our reverse strategy ranked first in the blue pool in the first live competition. It earned around \$24,800 during the competition. Figure 3 shows the value and holdings of the reverse strategy and the price chart of MSFT in real time during the competition.

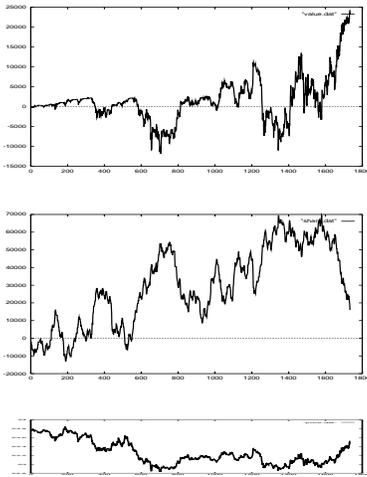


Figure 3: PLAT First Live-data Competition, Dec. 18, 2002

After this initial competition, one additional informal live competition was held on the following day. The participants were our team and the three top scoring teams from the original red pool. Our reverse strategy repeated as the winning strategy on that day. Figure 4 shows the performance of our reverse strategy on that day. Since the competitions were only on two days, the results are not be taken as conclusive evidence of our strategy's effectiveness.

The second PLAT live competition was held over 15 trading days from February 24th to March 18th 2003. It averaged the performance over many simulations to get more reliable indications of the performances of the trading strategies. The competition was structured as 3-week, 3-round tournament with one week (5 trading days) for each round. In the

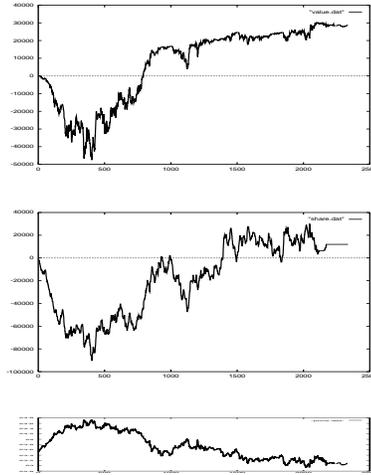


Figure 4: PLAT Informal Live-data Competition, Dec. 19, 2002

first round, 13 teams (two new teams participated in the second competition) were divided into two pools. The total returns for the first round were aggregated and three top scoring teams from each pool advanced to the second round. In the second round, 6 remaining teams were run in a single simulation. The top two performing teams in the second round met in the finals. Again, Each agent was restricted to maintain holdings within 100,000 of 0 (positive or negative). Violations were grounds for disqualification from the competition. Our agent was assigned in the blue pool in the first round, and was Team 5. Table 9 and Table 10 show the results from the first round. Team 11 in the red pool was disqualified after three trading days because of repeated and extreme violation of the share position limit. Team 7 had third-place earnings in the red pool, but was disqualified for its share position limit violations three times in five days. The red pool qualifiers were Teams 12, 8 and 13. The blue pool qualifiers were Teams 4, 5(our trading agent), and 1.

Team	24-Feb	26-Feb	27-Feb	28-Feb	3-Mar	total	Rank
Team 1	-479	3711	17603	1278	-266	21847	3
Team 2	334	2841	-53	-464	338	2996	5
Team 3	1046	3845	1980	379	1644	8894	4
Team 4	100557	-3104	-7314	-1642	-4360	84137	1
Team 5	12489	26357	-4304	12236	-346	46432	2
Team 6	-15228	-79442	-2052	4218	-20816	-113370	6

Table 9: PLAT Second Live-data Competition Round 1, Blue Pool

The six qualifiers advanced to the second round, facing off in a single pool over a five-day trading period to determine the two finalists to meet in the third round. The second round ran from March 5th, 2003 to March 11th, 2003. Table 12 shows the results of the second round.

DATE	Ticks	$\bar{\delta}(s.t.d.)$	Slope a	# shares	Unmat ched	L.M.	P.L.M.	E.C.F.	A.P.T.	Value
Dec05, 2002	4342	.0176(.0145)	-.000447	2,517,000	74,369(.030)	23,975	23,156	21,032	22,767	-43,583
Dec06, 2002	4464	.0186(.0151)	.000129	2,612,000	64,522(.025)	23,985	23,910	21,594	12,722	17,313
Dec09, 2002	4509	.0178(.0144)	-.000373	2,392,000	79,753(.033)	22,232	22,084	18,102	16,327	-30,257
Dec10, 2002	4663	.0163(.0137)	.000086	2,368,000	39,517(.017)	19,067	18,895	19,293	16,815	31,347
Dec11, 2002	4200	.0149(.0118)	.000248	2,380,000	63,816(.027)	17,449	17,365	7,837	-11,101	-21,637
Dec12, 2002	4719	.0149(.0126)	-.000191	2,342,000	60,647(.026)	17,836	17,960	10,743	18,056	17,006
Dec13, 2002	4656	.0142(.0130)	-.000249	2,261,000	115,941(.051)	16,961	16,917	-10,948	-24,488	-48,014
Dec17, 2002	4700	.0134(.0115)	-.000062	2,158,000	37,012(.017)	14,699	14,549	14,391	23,741	17,653
Dec18, 2002	4312	.0163(.0133)	-.000026	2,509,000	73,943(.030)	20,608	20,448	5,509	29,607	2,191
Dec20, 2002	4550	.0140(.0142)	-.000084	2,172,000	26,005(.012)	15,575	15,563	-51,828	-8,884	5,600
Dec23, 2002	4645	.0135(.0122)	.000166	1,910,000	56,545(.030)	13,244	13,107	-72,643	-35,129	-40,556
Dec27, 2002	4404	.0128(.0109)	-.000093	2,131,000	70,031(.033)	13,846	13,746	13,222	16,056	-2,174

Table 11: Historical Experiment Using the Reverse Strategy

Team	24-Feb	26-Feb	27-Feb	28-Feb	3-Mar	total	Rank
Team 7	-23935	20338	20791	50949	-28460	39683	3
Team 8	-9590	8200	79008	21357	-20795	78180	2
Team 9	2243	4414	6038	2103	3072	17870	6
Team 10	532	-2915	38396	-549	-3088	32376	5
Team 11	-4141632	-3434560	-31561216	N/A	N/A	N/A	7
Team 12	20187	29090	54001	7208	16673	127159	1
Team 13	-15066	8572	42734	25170	-26168	35242	4

Table 10: PLAT Second Live-data Competition Round 1, Red Pool

Team	5-Mar	6-Mar	7-Mar	10-Mar	11-Mar	total	Rank
Team 8	22433	18640	-35475	-9826	-4156	-8384	4
Team 12	-14840	27982	-28438	13257	23120	21081	1
Team 13	-4223	17731	-40912	-18271	21933	-23742	5
Team 1	-563	-395	-967	-776	-486	-3187	3
Team 4	-6700	6	-30608	83	-135	-37354	6
Team 5	11095	12105	-21931	10399	-9979	1689	2

Table 12: PLAT Second Live-data Competition Round 2

The two top scoring teams (Teams 12 and 5) from the second round advanced to the finals. They were the only two teams in the black after the second round competition. Table 13 shows the result of the third round competition.

Team	12-Mar	13-Mar	14-Mar	17-Mar	18-Mar	total	Rank
Team 12	-54156	-79233	-21896	-18032	-886	-174203	2
Team 5	-33382	-54611	6915	-71397	26304	-126171	1

Table 13: PLAT Second Live-data Competition Round 3

Despite big losses in the third round, our trading agent lost less money than our competitor, and we (Team 5) won the competition.

Figure 5 and Figure 6 show the performance of the reverse strategy on two trading days during the finals. First, on Mar. 17, 2003, Wall Street had its best day since Oct. 2002. However, it was a disaster for our reverse strategy. The MSFT stock price increased consistently on that day, which, as per our analysis in Section 4, is the worst case scenario for the reverse strategy. As a result, we lost a lot of (simulated!) money. In contrast, on Mar. 18, 2003, The MSFT price roughly oscillated around a constant line, and

our reverse strategy finished in the black.

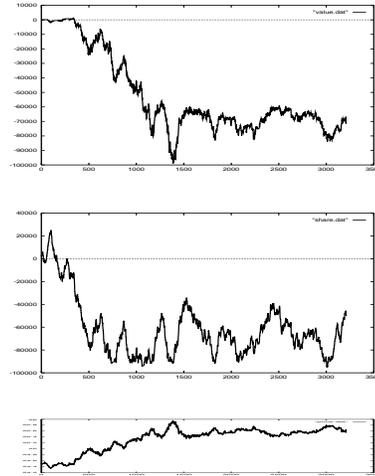


Figure 5: PLAT Second Live-data Competition Round 3, Mar. 17, 2003

Besides the live-data competitions, we also did some historical experiments with the reverse strategy. Table 11 shows the results. The dates in the table are not continuous because the NASDAQ is closed on weekends and because there were some network problems preventing the PLAT simulator from getting data from the ECNs on certain days. Column 2 shows the number of simulator ticks in the day. For the purposes of aligning with our models, we assume every four ticks represents one cycle. Column 3 shows the average microvariation. The standard deviation is listed in parentheses. Since the real world data is very noisy, the standard deviation value is large. Column 4 is the slope of the price trend, which is determined by the difference between the closing price and the opening price divided by the number of ticks in the trading day. Column 5 is the total number of shares in buy orders and sell orders placed by the reverse

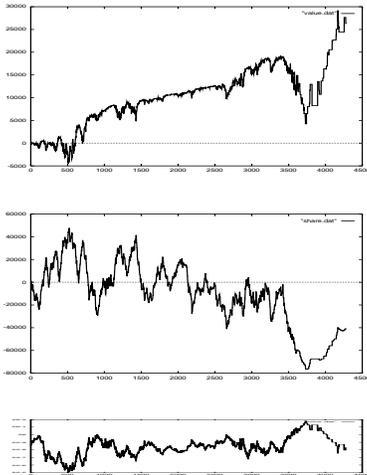


Figure 6: PLAT Second Live-data Competition Round 3, Mar. 18, 2003

strategy within that day. Column 6 is the number of shares left in the queues after the trading session is closed. The unmatched rate of shares is listed in parentheses. L.M. is the calculated value with the linear model. P.L.M. is the calculated value with the piece-wise linear model in which we manually break the actual price trajectories into coherent line segments and compute a different δ for each line segment. Due to the nature of the calculation, these values are very close to those in the linear model. E.C.F. is the calculated value with extended-cycle-frequency model. The next column shows the value assuming full matching, but following the actual price trajectory. Finally, we list the actual value earned by the reverse strategy.

From Table 11, we can see that if the unmatched rate is high ($\geq .03$), the reverse strategy loses a lot of money. If the reverse strategy has too many unmatched orders, it indicates that the price trend is biased. In an extended upward trend, the reverse strategy will have many sell orders left in the queue, while in an extended downward trend, it will leave many orders in the buy queue. As shown in our analysis, if the price trend is biased, the reverse strategy tends to lose money quadratically.

The linear model is an optimistic model. In the linear model, we assume that the price trend is not biased and the price oscillates around a straight line. It is so optimistic that the value we have in the linear model is the upper bound that the reverse strategy can make for most of the trading days.

The values of the extended-cycle-frequency model are quite close to the values in the fully-traded model. The only difference between these two models is whether they have different microvariations or not. When the unmatched rate

is low, the value of these two models is quite close to the value from the actual simulation; otherwise, there are big differences. All in all, it is clear that none of our models fit real market conditions consistently closely. Nonetheless, they help us gain an understanding of the conditions under which the reverse strategy is expected to succeed.

6. FUTURE WORK

As is apparent from the PLAT live-data competition and Table 11, under some market conditions, the reverse strategy loses a lot of money. We have postulated that these losses are a result of the price trend going in one direction consistently. In such market conditions, the basic strategy, which does the opposite of the reverse, will make profits. Thus, if the bias of the price trend can be detected, we can use the basic strategy, instead of the reverse strategy. If we can synthesize the basic strategy and the reverse strategy in the simulation, we hope the trading strategy will be more profitable.

In addition, the strategies presented in this paper ignore the most unique feature of the PLAT simulator, namely the ability to use complete order book information to influence bidding decisions. Though we are encouraged by the performance of the reverse strategy given its simplicity, we hope to improve its performance by incorporating order book information, perhaps to predict the appropriate times to switch between the basic and reverse strategies as mentioned above.

7. ACKNOWLEDGMENTS

We thank Professor Michael Kearns and his PLAT group at the University of Pennsylvania for their development of the PLAT simulator and for allowing us to participate in the competitions.

8. REFERENCES

- [1] The Penn-Lehman Automated Trading Project: <http://www.cis.upenn.edu/~mkearns/projects/pat.html>
- [2] Archipelago: <http://www.tradearea.com>
- [3] Bloomberg: <http://www.bloomberg.com>
- [4] Island: <http://www.island.com>
- [5] The 2002 annual report of Securities Industry Foundation for Economic Education, 2002: http://www.sia.com/about_sia/pdf/annual2002.pdf
- [6] Virtual Stock Exchange: <http://www.virtualstockexchange.com>
- [7] R. Freedman et. al., eds., *Artificial Intelligence in Capital Markets*, Chicago, IL: Probus Pub., 1995.
- [8] A. Skabar and I. Cloete. "Discovery of Financial Trading Rules." In *Proc. Artificial Intelligence and Applications*, pp. 121-125, 2001.
- [9] J. Kingdon. *Intelligent Systems and Financial Forecasting*. New York, NY: Springer, 1997.
- [10] M. Sheimo. *Stock Market Rules: 50 of the most widely held investment axioms explained, examined and exposed*, Chicago, IL: Probus Pub., 1991.
- [11] P. Stone and A. Greenwald, "The First International Trading Agent Competition: Autonomous Bidding Agents". *Electronic Commerce Research*, 2003 (to appear).