

Penn-Lehman Automated Trading Project

Anant Jalan

Email: ajalan@seas.upenn.edu

Icksoo Jung

Email: ijung@seas.upenn.edu

Faculty Advisor:

Michael Kearns

Email: mkearns@cis.upenn.edu

ABSTRACT:

PLAT (Penn-Lehman Automated Trading) Project is a broad investigation of algorithms and strategies for automated trading in financial markets. The PLAT Project's centerpiece is the Penn Exchange Simulator (PXS), a software simulator for automated stock trading that merges automated client orders for shares with real-world, real-time order data. The following strategies are frequently used today in the market: Channel Breakouts, Relative Strength Index, Moving Average Crossovers, Momentum Strategies and Market-Making. The existing strategies perform well in very specialized market conditions and badly in all other conditions. Our goal was to use a combination of these strategies to develop our own automated trading agent that would be successful in a broader variety of market environments. These strategies are to be tested through end of semester competitions.

BRIEF DESCRIPTION OF EXISTING STRATEGIES:

We implemented the following strategies and observed how they did when run against each other and a Background agent (uses some type of a normal distribution to place random buy/sell orders) when run on the PXS.

Our analyses of the results yielded the following useful insights:

Market Making (MM) strategy:

This is clearly the most dominant strategy in our simulations. Clearly, the market making strategy is successful in a volatile market with spikes and troughs in short time intervals. This is the only strategy that does not bet on a trend. The others are betting on the fact that if the price has started to rise, they will rise more and vice versa. The market making is betting on the fact that there will be great short-run volatility in the market.

What market making does is that it places a buy order just below the market's bid and a sell order just above the market's ask. This is updated every time interval according to the movement in the market's bid and ask. The way this strategy makes money is as follows:

If the price first rises/falls so that the market's ask/bid goes above/below the market maker's sell/buy, the market maker trades and sells/buys a fixed amount of stock. Until then, the market maker does not trade. In the same time interval, if the market price now falls/rises so that the market's bid/ask goes below/above the market maker's buy/sell orders, the market maker now trades the same amount of stock in the opposite direction (i.e. buys/sells). As a result, in that time interval, it has no net position in the stock. However, it has succeeded in making the spread in its buy and sell orders. For example, if market's ask/bid is 23.2/23.1, and the market maker places a sell/buy order at 23.3/23.0, and what we mentioned above happens, the market maker makes $23.3 - 23.0 = 0.3$. This is how it makes money. It makes money when the price moves up and down in the same time interval (quickly). So it thrives in a volatile market.

Looking at the flip side, if the market follows a trend and there is no short run volatility, this strategy loses money. This happens due to the following reason: If the market price falls (bid and ask go down), the market maker starts to buy (as his buy is now higher than market's

bid). Now if there is no reversal in the same period, the market maker changes his buy/sell prices around the market's new bid/ask. If the same thing happens the next period, the market maker again buys some stock. So in effect, the market maker is buying at high prices (relative to market) when the price is falling. So he starts to lose money. If the trend continues, he loses more and more. When price does finally start to rise (if it does before liquidation), the market maker starts selling at a price lower than market. So he loses money.

By looking at the output from some of the simulations in which MM participated, it makes a lot of money when run against any of the first four strategies. When the other strategies are run without MM, the markets are more volatile in the first half of the day and less volatile (smooth downward trend) in the second half of the day. However, when MM is part of the simulation, the volatility continues throughout the day. This does not happen when MM is run against MM. So only when MM runs against any of the other 4 strategies, volatility is created. We speculated this happens because: When price is falling, the other 4 strategies are going to start selling and market making will start buying. This will move stock price back up and then the other 4 strategies will start buying and market making will start selling. This keeps happening and volatility is created. Market making makes the spread each time and the other four strategies lose each time. When MM is run with MM, there is no one to sell when prices are falling and so there is less volatility in the second half of the day.

Note: Market making is the only strategy that is a price-setter (places limit orders). The other strategies are price-takers (place market orders).

Momentum (MO) Strategy:

This is one of the trend strategies. It loses money when there is volatility in the market (spikes). It makes money if there is a trend and the trend turns around at some point (or there is liquidation).

Clearly volatility is bad for the strategy. If there are lots of spikes, just like the other trend strategies, MO will lose the bid/ask spread. Let us assume the number of intervals in consideration is 20. If the price goes up in first 20, buy is placed. If it then goes down in the next 20, a sell order is placed. So the strategy is buying at high prices and selling at low prices in the case of a volatile market.

If there is a trend, this strategy may make money. Suppose there is a steep up-trend, MO will buy some stock. Now let's say the up trend continues, however less steep (less than the 0.1 threshold in the required 20 interval time frame). So now there is no new buy in but the initial shares bought are appreciating (as stock price still rising but slowly). Now after a long time if there is a decrease in stock price, the strategy sells its shares and can make a profit. Also, if there is a forced liquidation at day end and the BG is set to place a high volume of orders (such that the sudden liquidation does not move stock price too much), the strategy will make money. Therefore, a better way to implement the strategy would be allowing the agent to liquidate all of its position when it predicts the downward movement. The same argument applies for a constant down trend.

In conclusion, this strategy bets on the fact that if the price is going up /down, it will keep doing so for a while. So it bets that a trend will be maintained. As a result, lack of trend and short run volatility kills this strategy.

Moving Average (MA):

Moving Average is a trading strategy that bets on the trend of stock price movement. The trade is triggered when the current price line crosses the average price over a time period specified as an argument. If current price up-crosses the moving average, it will take a long position on the stock, betting on the continuous upward movement. Conversely, if current price down-crosses the moving average, it will take a short position betting on continuous downward movement.

Throughout the simulations, Moving Average Agent was consistently losing money. This is due to high volatility in stock price over short run period. The agent will buy or sell the shares through market order whenever the price crosses the average line. This will result in loss of Bid-Ask spread, which is not compensated by the profit from a long run trend. The agent will not take two consecutive long or short positions in stock since the current price can only up or down cross the average line once before it reverts its position. This is the main difference between Moving Average strategy and Momentum strategy.

For this strategy to be profitable, two conditions have to be met. First, the volatility of stock over a short run period should not be high, so the agent would not continuously trade. Second, the long run trend should exist, since the agent is betting on the assumption that crossing the average line is a signal of upward or downward trend in stock price movement. In absence of high short-term volatility, the agent will make money at the inflection point of the long-term trend.

Channel Breakout (CB):

Channel Breakout is a strategy that also bets on the trend of stock price movement. The agent records the highest and lowest price over a period of time specified as an argument and trades if current price is out of the range of the channel. In case of upward breakout, it will place a market buy order. If the current price is lower than the lower bound of the channel, it will place a market sell order. Channel Breakout is a more conservative trading strategy in a volatile market than Momentum Strategy. Since the agent examines a period of time rather than a point in time to make a decision, it will give a bigger range for no trade, which will make losses from bid ask spread smaller. However, unlike Moving Average strategy, this strategy can take consecutive long or short positions, which can magnify its return in an upward or downward trend.

Throughout the simulations, Channel Breakout lost money except when it was trading against another Channel Breakout Agent. For this strategy to be profitable, similar conditions as MA strategy have to be met. There has to exist a long run trend in stock price movement and the volatility in short run deteriorates the return of this strategy since it can recognize if the movement is a long run trend or a short run volatility.

Relative Strength Index (RSI):

Relative Strength Index (RSI) is trading agent that believes that the RSI, which is computed over a time period, can predict the long run trend of the price of stock. RSI is calculated by

$$RSI = 100 - \frac{100}{1 + \frac{U}{D}},$$
 where U is the average of up price movement and D is the average of

down price movement. Given lower and upper thresholds, the agent buys when the current RSI up-crosses the lower threshold and sells if the current RSI down-crosses the upper threshold. This allows the agent to take consecutive long positions and the strategy will not change its view on the stock movement as often. Short-term volatility of the stock price doesn't affect the position of the agents, which prevents it from losing money on spread.

RSI performed relatively well over numerous simulations since it did not lose money over short-term volatility. For this strategy to make money, the predicted trend by RSI should be correct. Since volatility doesn't play an important role in this strategy, as long as the long-term trend is correctly predicted, the agent will be profitable.

DECEMBER 2004 COMPETITION:

Description:

An end of semester competition was held last semester in December 2004. The competition comprised of a set of different simulations. Each simulation was run with different parameters so that the effectiveness of a strategy could be tested over a wider variety of market conditions. For example, some simulations produced trending markets, some produced volatile markets and some produced inactive markets. Besides, the available liquidity was varied in each simulation. In order to achieve such conditions, the competing strategies were run along with a mixture of different background agents and other existing strategies. The basic idea was to produce a competition environment where different things happened on different days and no agent knows before hand what exactly is going to happen on any given day. This was done intentionally to try and make this simulation as "real world" as possible.

Evaluation Criteria:

(a) In terms of raw performance, there will be a single criterion for the competition, which is the Sharpe ratio of each client's daily profit and loss. More precisely, suppose that over x trading days, the final profits or losses of each client (see below for how this is computed) are p1, p2, p3,..., px, where each value might be positive (profit) or negative (loss). Then each client will be judged by the value (average of p1, p2,...,px)/(standard deviation of p1,p2,...,px) which is (one form of) the Sharpe ratio.

(b) There are no limits on how many shares each client can buy or sell in a day, but each client must completely liquidate its position (act like a day trader) by the close of each trading

day. There will be a monetary penalty applied to the client's profit or loss each day according to how many shares it fails to liquidate by the close. Each share that a client has a long position at the close will be valued at 0. Each share that a client has a short position in at the close will be valued at twice the closing price, and the client's profit/loss will be docked accordingly.

(c) PXS will be run in transaction cost/rebate mode for the competition. Every time PXS executes a trade, one side of the order must have already been sitting in one of the order book, and the other side of the order must have been the "incoming" order. For each share executed by PXS, the party whose order was already in the books shall receive a Rebate of \$0.002, and the party that was the incoming order shall pay a transaction Fee of \$0.003. This is exactly the policy used in the "real world".

Our client's old strategy:

Since Market-Making was the dominant strategy based on our simulations, we decided to base our strategy off MM. As discussed, the Market-Making agent makes a profit in most market conditions except when there is not enough up-down volatility.

Since volatility was the most important factor for the success of this strategy, we decided to compute a volatility index based on recent prices. We recorded the maximum and minimum price reached in every interval. If the average difference between the max and min prices over any period was greater than half the size of delta, we employed our modified Market-Making strategy. Our modified MM strategy was similar to the generic Market-Making strategy. However, we made two modifications: asymmetric orders and liquidity-adjusted-volumes. We placed asymmetrical buy/sell orders i.e. we did not use the same delta for bids and asks. Hence, if there was a weak up trend, our ask delta was higher than our bid delta as we expected prices to rise further. Just the opposite was done for weak down trend markets. That is, rather than placing orders symmetrically around the current price, the agent places buy orders very close to the current price and places sell orders at higher prices in a weak up-trend market and the reverse in a weak down-trend market. For liquidity-adjusted-volume, we adjusted the number of shares per order according to the traded volume of the stock. If the volume traded was high, the market was liquid and we were more confident taking larger positions. If volumes traded were thin, we placed smaller orders.

In addition, we recorded the price at the end of each interval. We used these prices to calculate a trend index. If the prices at the end of the last few intervals had been increasing steadily, then the market was seeing a trend. When the net price movement over the last five minutes was greater than twice the average volatility over that period, we turn off our modified MM strategy and turn on our trend strategy. The trend strategy simply buys at the market price in the case of an up-trend and sells at the market price in case of a downtrend.

Results:

The following table summarizes our results over the numerous trials. The first column indicated our results before the application of liquidation penalties and the second indicates the results after the application of relevant liquidation penalties.

	Before Liquidation Penalty	After Liquidation Penalty
SBG+CB+X	625.0	625.0
SBG+MA+X	1,330.0	(3,967.0)
SBG+MM+X	1,353.0	(41,944.0)
SBG+MO+X	424.0	424.0
SBG+RSI+X	691.0	(173,732.0)
SBG+X+X average	690.0	690.0
SBG+ALL+X	257.0	257.0
SBG+X1...X8	(9.0)	(9.0)
ABGup+ALL+X	(1,455.0)	(1,455.0)
ABGdown+ALL+X	(2,432.0)	(2,432.0)
REALBG_dell+X	8,645.0	(1,548,461.0)
REALBG_msft+X	17,857.0	(4,139.0)
REALBG_yhoo+X	12,473.0	(103,562.0)
Average	3,111.46	(144,438.85)
Standard Deviation	6,031.51	425,169.29
Sharpe Ratio	0.516	(0.340)
Rank	2	4

Analysis of our performance:

As seen, we were placed second before liquidation penalties and fourth after liquidation penalties. In most simulations we were not able to liquidate fully and were really hurt by the liquidation penalties. Since we did not limit our position, we took on positions of very large sizes and we were unable to unload these large positions by the close of the day. Given that most agents probably trade using similar strategies, they all take similar positions (either all long or all short). At the end of the day, all these agents must liquidate. Since they are either all long or all short, they scramble for the same trades i.e. either everyone is selling or everyone is buying. This means it is harder for our agent's trades to go through and we are left with a substantial position size at the end of the day.

Additionally, the goal of the competition was to maximize the Sharpe Ratio (average/standard deviation) or our profits and not simply make the most profit. Hence, it was more important to be consistently profitable rather than very profitable. Our results were not very consistent across different simulations. We either made a lot of money or lost a lot of money. The adjustment of our trading volume based on volume traded was probably not a good idea.

Furthermore, our measure of volatility was not effective. We needed to measure the volatility to predict if the price would cross both the bid and ask orders places by us. However, recording the difference between max price and min price did not provide us with sufficient information to make this prediction. We could not distinguish if the volatility was around the current price or solely due to upward or downward movements. We may need to revise the estimates for our delta.

There was something wrong with our transition to the trend strategy. In the trials with the asymmetrical background agent (trending markets), we perform poorly. We probably failed to turn on our trend-based strategy and continued to trade with our volatility-based strategy.

The primary weakness of our strategy was the inability to liquidate completely.

MAY 2005 COMPETITION:

There will be another competition held sometime in May 2005. The nature of this competition will be identical to that of the December 2004 competition and the evaluation criteria will be the same. Having learnt from the results of the December, 2004 competition, we revamped our agent and incorporates several strategic changes.

Our client's new strategy:

Given the results, there were several issues we needed to address:

1. **Liquidation:** We were not liquidating in all circumstances. The liquidation penalty is extremely severe and making sure we liquidate should be our first priority.
2. **Variance:** Since the evaluation criteria is Sharpe Ratio and not average profits, it is important to be consistently profitable rather than just very profitable.
3. **Trend Strategy:** In the December competition, we did well in volatile markets (SBG). However, in the case of trending markets (ABG), we performed badly and lost money even before liquidation penalties.

Maximum Position Bound Strategy:

In order to address the first two issues, we implemented the following system. At any given instant we maintain a maximum bound for our share position (both long and short). This bound is set to 1000 at the beginning of the day. Every interval of time, this bound is reduced. The idea is that this should converge to zero by the end of the day. Earlier on in the day, we reduce bounds after longer intervals (every one hour). As the day moves on, we cut the maximum size of our bound more frequently.

If our share position ever exceeds this bound, our agent acts as follows:

If we are long too many shares, our agent withdraws all the buy orders in the book. Additionally it places a sell order at the current market prices minus a desperation value. The idea is to reduce the long position immediately. If the order is not executed and we are still above the bound, we increase desperation value and place a sell order at a lower price. We keep doing this until we are able to sell.

If we are short too many shares, our agent does the exact opposite. It withdraws all the sell orders from the book and places a buy order at the current price plus the desperation index and keeps doing so until the position comes back within the bounds.

So each time we violate a bound, we do not do any active trading. Instead we concentrate on reducing position size until we are within the bound. Once we are within the bound, active trading begins again. The desperation value is incremented by half a cent each time.

This serves two purposes. First it helps us limit our position and helps prevent liquidation problems later on in the day. Second, by keeping a standard maximum position size, we are reducing the variance in the agent's profits over simulations.

Momentum Based Strategy:

In order to address the problem with the trend strategy, we revamped our entire momentum based trading strategy. Our new strategy works as follows:

Each trading interval is 10 seconds long. We record the price at the end of the last 20 trading intervals. At each interval we compare the current price to the price recorded 20 intervals back i.e. we compare it to the price 200 seconds ago. If the price has risen by more than 1 cent, we increment a counter. We keep incrementing this counter every interval as long as the price in the current period is 1 cent higher than the price recorded 200 seconds before that period. If this counter reaches 30, we turn off our Market-Making strategy and turn on our momentum-based strategy. We also withdraw any orders placed by Market-Making that are still in the book. Then we buy at the current price in accordance with the trend. The volume that we buy is either half of the maximum position or whatever we can buy without exceeding the maximum position, whichever is smaller. We stay in the trend mode until we see an interval where we see a decrease of price from the price 200 seconds ago. If we see this decrease, we turn off the trend strategy, turn back on the Market-Making strategy and reset the counter to 0. This counter can start incrementing again and if it reaches 30, we switch back to trend based. Note this description is for an upward trend. The exact same opposite applies to a downward trend.

We see that we turn the trend strategy on if we have seen 30 periods of an increase of more than 1 cent from the price 200 seconds ago without seeing a fall in price. If we see 10 periods of increase of 1 cent and then a few periods of an increase of 0.1 cent and then 20 more periods of an increase of more than 1 cent, the market still has a trend and we turn on the trend strategy. Note the counter does not get reset to zero if there is an increase but not quite 1 cent. However, to increment the counter, we need to see a one-cent move.

We arrived at the parameters of 30 periods and 200 seconds after running various simulations with different parameters. We found these to be most optimal over various trials. Note also that we do not trade in the first 15 minutes of the day, as we require some time to detect a trend.

Market-Making Strategy:

Our Market-Making strategy runs when our trend strategy is shut off and vice versa. At one point, only one is running. If there is no trend, Market-Making is effective and should be turned on. We use a very basic Market-Making strategy with a constant delta and a constant volume. We realized that varying volume does not help consistency. Hence at each interval,

we place buy orders at delta above the best ask and at delta below the best bid and wait for these orders to be hit.

We use a delta value of 0.008. We ran simulations with several delta values and found this value to be the most effective. The volume traded is just 10% of the maximum allowed position for that period.

Desperate Liquidation in the last 20 minutes:

In the last 20 minutes, we do not do any active trading. We simply attempt to liquidate our position. Since we have used converging bounds, we already expect to have a very small position by now (about 5 shares). We withdraw all orders currently in the book. If we are long a certain number of shares, we place a sell order for 1 cent. If we are short a certain number of shares, we place a buy order at 2 times the current price minus 1 cent (as this is what the penalty is for holding the share). Once we liquidate, we do nothing and wait for the trading session to close.

Results from sample test runs:

In preparation for the upcoming competition, we tested our agent in a wide-range of markets. We wanted to test the robustness of our strategy in varying market conditions. WE ran four trials. Each of these trials included our strategy and a mixture of other technical strategies such as Market-Making, Momentum, Relative Strength Index, Moving Average and Channel Breakout. In order to create trending markets, volatile markets and markets based on real world data, we used a mixture of background agents in each trial. We adjusted volumes traded by the background agents and the technical strategies so as to try and get markets with surplus liquidity and liquidity crunch. The results from each of these trails are given below:

Trial 1: We tried to create a smooth up-trending market with no short-term volatility. Hence we used an asymmetrical background agent with a big asymmetry in the buy/sell distributions. We included other technical strategies such as MM, RSI and MO in the same trial. *Please refer to appendix A for results.*

Trial 2: We tried to create an up-trending market with short-term volatility. Hence we used an asymmetrical background agent with a big asymmetry in the buy/sell distributions combined with a symmetrical background agent trading smaller volumes but at more frequent intervals. We included other technical strategies such as MM, RSI and MO in the same trial. *Please refer to appendix B for results.*

Trial 3: We tried to create a volatile market with no apparent trend. Hence we used a symmetrical background agent. We included other technical strategies such as MM, RSI and MO in the same trial. *Please refer to appendix C for results.*

Trial 4: We tried to create a real world market. Hence we used the real background agent of Yahoo stock on 2nd June. We included other technical strategies such as MM, RSI and MO in the same trial. *Please refer to appendix D for results.*

Trial	Cash at end	Shares at end	Rebate Shares	Fee Shares	Net Profit
1	\$919.01	0	861	2301	\$913.83
2	\$1,214.50	0	126201	5717	\$1,449.75
3	\$114.79	0	29917	3939	\$162.81
4	\$43.42	0	15954	2112	\$68.99
Average Profit		648.85			
Standard Deviation of Profits		654.25			
Sharpe Ratio		0.9917			

As seen from the results, our strategy was able to make a positive profit in each of the simulations. More importantly, we were able to liquidate our position completely in each of the trials. As seen from the general shape of our share position graph, we start out with amplitude of 2000 (1000 in each direction) and converge to zero as the day proceeds. This was what we wanted to see while implementing our strategy.

CONCLUSION:

Having learnt from the results of the December competition we were able to make some major changes to our client's trading strategy. We continued to use a basic strategy of Market-Making (for volatile markets) mixed with momentum. However, in order to ensure complete liquidation we kept strict bounds on our maximum share position. These bounds converged to zero as the trading session neared its end. On running sample simulations, our client's strategy performed well in a variety of market conditions. We are hopeful of a much better result in the May 2005 competition.

APPENDIX:

Appendix A – Trial 1 Results (Smooth up trend with no short term volatility)




Figure 1: Simulator values of the shares traded for the day




Figure 2: Shares Position of the client through the day




Figure 3: Simulator price for the day

Appendix B – Trial 2 Results (Up trend with short term volatility)




Figure 1: Simulator values of the shares traded for the day




Figure 2: Shares Position of the client through the day




Figure 3: Simulator price for the day

Appendix C- Trial 3 Results (High volatility with no real trend)




Figure 1: Simulator values of the shares traded for the day




Figure 2: Shares Position of the client through the day




Figure 3: Simulator price for the day

Appendix D– Trial 4 Results (Real Background Agent on Yahoo Stock)




Figure 1: Simulator values of the shares traded for the day

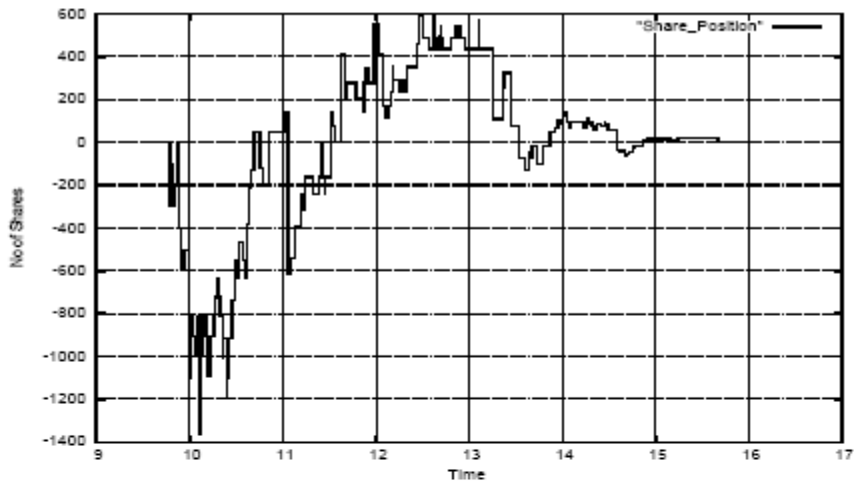


Figure 2: Shares Position of the client through the day




Figure 3: Simulator price for the day