Detecting Trading Patterns via Markov Decision Processes for Market Surveillance

Asoka Korale[1], Fuard Ahamed[2], Kaushalya Kularatnam[3], Liam Smith[4]

[1]asokak@millenniumit.com, [2]fuard@millenniumit.com, [3]KKularatnam@lseg.com,[4]LSmith2@lseg.com

Abstract—Investor confidence in the markets rests in the ability to detect illegal trading activities and manipulations carried out with the intention of influencing prices. Market regulators deploy intelligent algorithms to detect these activities and identify the individuals responsible to ensure a level playing field and fairness to all participants as illegal activities result in inefficiencies and higher costs.

In this paper we present an algorithm based on AI and machine learning techniques that estimates the average trading strategy of a trader by modeling the transactions performed in response to the observed state of the market and the expected profits and loses made with respect to each transaction. Through this modeling we can compare between the strategies of different traders in addition to capturing the actions of individual traders in response to market conditions. Through this we aim to infer activities that provide certain participants and unfair advantage over others, allowing us to learn newer ways of market manipulation.

While market prices are determined through a complex stochastic process underlying trading activities the factors that largely influence prices are observable through the state of the order book which represents the supply and demand for a stock at any given time. The state of the order book can be quantified in relative terms via attributes such as imbalances in volumes, spread in the bid & ask prices, liquidity and volatility of the stock and price depth of the book. This state then influences the actions which are the particular types of orders placed by traders in search of profit. For our initial analysis we will focus on limit order books that employ a price-time priority. This is due to the fact that this is the most common priority mechanism currently used. We relate each such order to a finite set of actions using the order book at a given point in time to determine its relative standing.

We cast this problem as a Markov decision process where each state is a multidimensional variable defined with respect to attributes that reflect the relative position of the order book, the reward returned to the trader by the profit or changes to the dynamics of the order book when an order is placed by a trader.

Trading strategies are in effect sequences of orders amenable for modeling via Markov decision processes that is time varying. Each subsequent state is dependent only on the current state and the action (order) and so follow a Markov process. The number of states is kept to a finite set by defining their characteristic attributes in relative terms. The actions too belong to a limited set as the types of orders are limited when defined in relative terms accounting for their particular position in the order book and relative volume.

The state transitions and model parameters are estimated by training on the actual states which are in effect a view of the price and order book that comes about in response to the interplay between orders placed by the trader and other influencing factors. In this way we generate an average view of the behavior of a trader based on his actions and state of the market.

Keywords—Market Survelillance; Markov Process; Markov decision process; Stochastic Processes; Trading Strategy

I. Introduction

Market regulators strive to ensure fair and proper functioning of the capital markets by enforcing legal provisions as a deterrent to market manipulation. Detecting market manipulations is important for maintaining investor confidence in the markets as manipulated markets are inefficient and bear additional costs to the investor. Regulators deploy surveillance algorithms to detect well-known types of manipulations that have patterns that are well defined. These patterns take the form of sequences of events that are defined in terms of certain types of orders, movements in the price and time intervals between such events.

Each order that is placed by a trader has a certain volume and price associated with it. The orders are placed in an order book in price-time priority which means that buy orders and sell orders are arranged in ascending and descending order of price in the buy and sell side of the order book respectively. The order book thus contains a measure of the demand and supply for a security as evidenced by the volume of orders at different price points placed on the buy and sell sides of the book. Typically it is the volume of orders at the top five or ten prices of the book that influence other traders to trade as the orders that are near the top of the book have the highest probability of executing.

The state of the order-book can be described in terms of the volumes on either side, depth in terms of the number of distinct price points in a side, the bid – ask spread which is the difference in top sell and buy prices among other attributes.

The strategy of a legitimate trader is then dependent on the state of the order book as it reflects the total demand and supply for a stock and the difference in buying and selling prices at any point in time. Therefore one would expect many traders to exhibit similar behaviors in their trading pattern depending on the strategy they follow in the long term as they are all trading with respect to a single source of demand and supply which they all observe on a single order book.

If one is able to gauge the strategy of individual traders it would be possible to compare their individual behaviors and determine those that are significantly different, especially before and after price sensitive news publications. This may aid compliment evidence leading to potential insider dealing related manipulations. Insider dealing is notoriously difficult to detect and prove due to the complex nature of relating market news with trading behavior. However the ability to detect trading behaviors that are markedly different to the rest of the trading community and then correlate those traders (behaviors) with news events that influence market prices will provide key insights in to potential illegal activity.

Since each order sent to the stock exchange is referenced with respect to a particular brokerage and a specific trader within each brokerage, it is possible to identify individual traders or groups of individual traders involved in whatever type of market activity they engage in. Thus collusive behavior which is the manipulative trading activity of several traders trading among themselves can also be identified if their trading patterns can be isolated.

II. Modelling the State of the Order Book

We devised an algorithm to analyse market data stored in a stock exchange and reconstruct the order book as a series of sequential events that reflect the series of actual orders coming in to the exchange for a particular security. Thus the program is aware of each order that is placed on either side of the order book at its respective price point. Thus the program has at its disposal complete knowledge of the order book at any given point in time. Using this knowledge two attributes that describe the state of the order book are defined.

Thus each new order or amendment to an existing order placed by a trader can be classified according to these two attributes which reflect the state of supply and demand in the order book at the time the order is placed. Through this we capture the high level trading strategy of a trader with each of his activities with respect to the relative state of the order book at a given point in time. In this modelling we consider four types of order events that accounts for more than 95% of all trading activity. They are the new order, cancel, amends and fills.

A. A measure of the price depth of an order

A vector is defined that captures the maximum, mean and minimum prices on a particular side of the order book. Using this vector the price of each new order is classified as falling in to one of two intervals namely, maximum to mean and mean to minimum.

B. A measure of the relative volume of an order

A vector is defined that captures the maximum, mean and minimum volume on a particular side of the order book. Using this vector the volume of each new order is classified as falling in to one of two intervals namely, maximum to mean and mean to minimum.

C. Relative state of an order

We classify each order placed by a trader in a way that it reflects its relative position in the order book at any given point in time. Its relative position is defined by comparing at the same time the volume of the order with respect to the distribution of the total volume in the order book and its price with respect to the set of unique prices in the order book. This results in a two dimensional state that is classified by a series of adaptive and time varying thresholds.

Table 1 gives the classification of each attribute of price and volume with respect to the number of unique prices in the order book and the distribution of order volumes in the order book respectively. Table 2 gives the classification of an order by price and volume.

TABLE I. CLASSIFICATION VIA DYNAMIC THRESHOLDS

Maximum - Mean	***Mean - Minimum***
High	Low

TABLE II. MULTIDIMENSIONAL STATE MAPPING

Volume		
Price	*High*	*Low*
High	1	2
Low	3	4

We also define an order cancel and an order-fill as states number 5 and 6. In this way we obtain six states each for buy and sell side for a total of twelve states when both buy and sell sides are considered together. Using these states we can cast the process of trading by a particular trader as the transitions between states of a finite state machine.

III. Markov Modelling of Trader Strategy

We model the orders placed by a trader as a sequence of events generated by a Markov random process where each event is mapped to one of the twelve states defined above. The transition matrix captures the conditional probability of transitioning between two successive states of this Markov model. The transition matrix is estimated from trading data for a particular security.

We divide the trading day in to a certain number consecutive trading events delimited by a window, so that each window encompasses a fixed number of events. We then observe the traders within each such window and model their trading events as a sequence of consecutive events. This enables us to compare the average trading behaviors between traders over similar time intervals. In this scheme each window will incorporate a different number of events for a given trader. The transition matrices however are conditional probabilities that are right stochastic matrices where all rows sum to one, which enables the comparison within and between traders.

A. Markov Processes

A first order Markov process is characterized by the memory less property so that in a sequence of random variables drawn from such a process, the distribution of the next random variable in the sequence is dependent only on the current one [1].

$$P(X_n / X_{n-1}X_{n-2}...X_1) = P(X_n / X_{n-1}) \qquad (1)$$

For the purposes of this paper, trading is modeled as being the result of a first order Markov process. However there is no restriction to higher order Markov models being considered and the analysis can proceed along similar lines by estimating the corresponding state transition matrices and anomaly detection via principal component & cluster analysis.

B. Stationarity

The probability distribution of a sequence of random variables drawn from a stationary a stationary distribution is invariant to a shift in time "*l*".

$$P(X_n X_{n-1} X_{n-2}...X_1) = P(X_{n+l} X_{n+l-1} X_{n+l-2}...X_{1+l}) \qquad (2)$$

The algorithm is adaptive in that it estimates a transition matrix for each delimited sequence of events. This transition matrix represents the dynamics in the underlying stochastic process which may not be stationary. The goal however is to compare between sequences and determine similarities and dissimilarities between sequences and in that respect the degree of stationary of the sequence will not pose a problem to the overall conclusions made with respect to detected outliers.

C. State transition matrices and trading strategy

Trading is essentially a sequence of events that are mapped to the states of a finite state machine modeled as a first order Markov process whose state transition probabilities are captured in a state transition matrix. The transition matrices capture the dynamics in the sequence and are thereby largely representative of the strategy employed.

As the transition matrices can be estimated over a series of windows defined over a series of consecutive events it is also possible to observe the dynamics of the strategy followed over time. Thus the changing strategies employed by a trader over time can also be quantified.

Alternatively one can calculate a state transition matrix for an extended period of time (like over a single days' worth of trading) for each trader and compare the average strategies employed across the traders.

IV. ANOMALY (OUTLIER) DETECTION CRITERIA

One of the main goals of this model is to identify traders who on average exhibit significantly different behaviors from the rest of the trading community. This is therefore an anomaly detection problem applied to the transition matrices which are representative of the average trading strategy followed. These are matrices of dimension 12 x 12 and thus have 144 elements when vectorized.

In one mode of operation each trader can have a sequence of such vectors for each transition matrix estimated over a series of consecutive windows. This type of analysis depicts the dynamics of the trading strategy of a trader over time. On the other hand one can estimate a single average transition matrix for each trader over an extended period of time. This type of analysis allows the comparison of each trader's strategy with the rest of the field. It is in this type of analysis the traders who's behaviors are markedly different from the rest can be identified. Once such anomalous traders are discovered they can be tagged for further investigation or monitored more closely for manipulations.

A. Principal components analysis in the outlier detection

The vectorized transition matrices have 144 dimensions and as such are difficult to analyze using traditional outlier detection methods like cluster analysis. Thus we employ principal component analysis (PCA) to reduce the dimensionality of the vectorized strategy matrices and compare between them.

Principal component analysis replaces the original variables with a reduced number of variables (dimensions) while accounting for as much variance in the original data set as is possible with the reduced set. All of the principal components are mutually orthogonal and represent decreasing amounts of total variation. Thus the first component accounts for the largest amount of variation, the second component the second largest amount of total variation and so on. As the PCA procedure replaces variables in the original data space with new variables (components) that are uncorrelated with each other, the vectors that are correlated in the original data space appear as groups of data points or clusters in the transformed space. Those vectors that are uncorrelated in the original data space appear as outliers in the principal component space.

Since our goal is to detect traders with widely different strategies from the norm, detecting correlation based outliers is a suitable strategy. Principal component analysis in this regard provides the necessary tools to discover those uncorrelated data points.

V. RESULTS

In section A we present results for the dynamic trading strategy of a single trader over an extended period of time like a single trading day. In section B we present a comparison of the average trading behavior across different traders.

A. *Dynamics of a trading strategy*

In this model the trading day was divided into windowed intervals of 5000 events and the state transition matrix for a single trader estimated over each window. In this particular case there were 236,000 trading events by all traders in the market for this security giving rise to 46 windows, each of size 5000 events. In this way the trading dynamics of a single trader is estimated over an entire trading day

Principal components analysis while regularly employed as a dimensionality reduction technique also finds application as an outlier detection method. The traders whose trading strategies that are correlated in the original space appear as clusters in the transformed principal components space while those traders whose strategies that are uncorrelated in the original space appear as outliers in the transformed space. When one uses the heuristic that outliers are data points that are those points that are most different from the rest of the data, the distance of the point from the origin provides an estimate for the degree to which it differs from the rest of the data.

The trading strategy is a multidimensional quantity due to its representation in a transition matrix estimated over a specific period of time. Thus it's highly suited for analysis via a multidimensional analysis technique such as principal components which is suitable for detecting correlations in higher dimensions.

The transition matrices so estimated for a particular trader was subjected to an analysis by the method of principal components where it was found that many vectors were largely correlated with each other allowing their representation in a reduced space of about 10 components. In Figure 1, we depict the first three principal components and observe the clustering of the data that is due to this correlation. We also observe a few outliers that are the transition matrices that are uncorrelated with the rest of the data.

Figure 2 presents the distance of the principal component scores from the orgin as a mesure of comparison between the transformed data points. Through this we observe that many of the data points are indeed clustered together in the principal componet space indicating their correlation in the originl data space. We conclude that over many windows the trading strategy of this trader has remaind largely similar. We can also use this analysis to detect those instances where the trader's strategy differs markedly from his own past behavior.

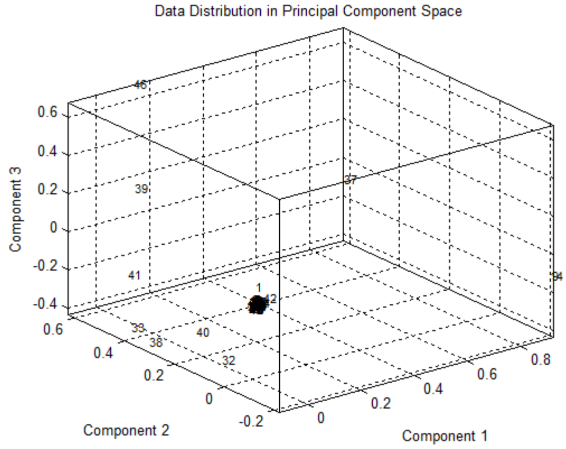

Fig. 1. Disribution of the original data in the principal component space

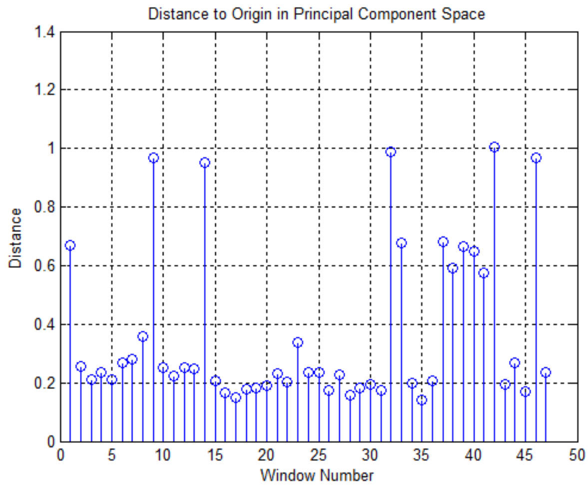

Fig. 2. Distnce of transformed data in principal component space

B. *Average behavior across traders*

The traders that had at least 2000 trading events across the day in question were selected for this analysis. Their individual strategies were captured via a single state transition matrix calculated on their trading pattern during the day. Figure 3 depicts the first three principal component scores of those 11 traders. We once again observe some clustering where the average behaviors (strategies) show some similarities.

The distance measure of figure 4 also confirms that several traders share broadly similar strategies while there are a few that are somewhat different. The transition matrices are somewhat dependent on the selected threshold setting scheme of which several have been explored. These include using non uniform intervals with tighter margins between top and next level, setting levels as an offset from the top of book price points or using the distribution of prices and volumes in their determination. The detected outliers are however consistent with a particular scheme as it is the relative behavior between traders that is being assessed.

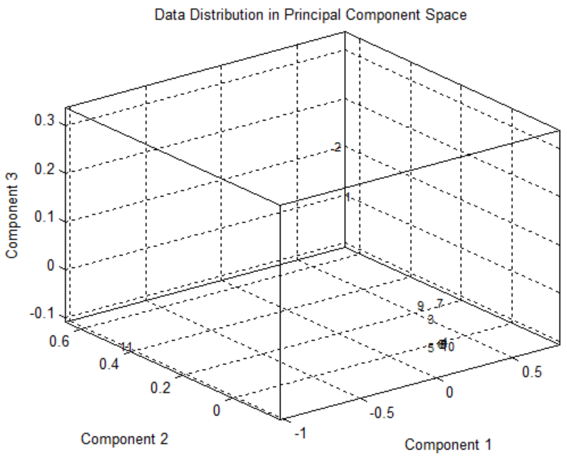

Fig. 3. Disribution of the original data in the principal component space

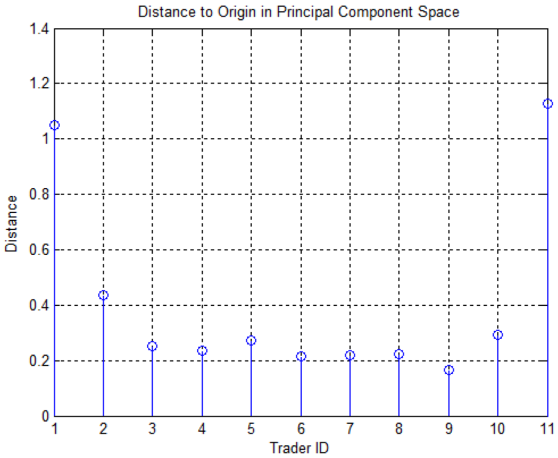

Fig. 4. Distance of transformed data in principal component space

VI. Conclusion

In this paper we developed a novel way of estimating the trading strategy of a trader by modeling the trading process as first order Markov process with multidimensional states that are determined via a set of dynamic thresholds. The states of the model are derived from the state of the order book and reflect the state of the market for a particular security at any given point in time.

We cast the orders submitted by a trader as a series of consecutive events that are mapped to a finite set of states in a way as to reflect the current supply and demand for that security. This was done by creating a software model of the order book to relate the position of the submitted order to its relative position in the order book.
The trading therefore is modeled via a finite state machine where certain states are defined via a set of dynamic or time varying thresholds that reflect the supply and demand for that stock at any given point in time.

As modeled, the multidimensional states based on the order book are sensitive to both price and volume variations across time and any deviation from the norm may readily be observed.

As the strategy of most traders depends on the state of the market as depicted in the order book, we observed some similarities in the overall average trading strategies across a group of traders on this day.

We also observed a relative similarity of the average trading strategy across different time periods for a given trader.

These broad conclusions with respect to the trading strategies of different traders are possible because we incorporate knowledge of the state of the market in the states of the finite machine.

Our modeling can also be used to detect anomalous trading behaviors by identifying those traders whose average strategies vary markedly from the rest of the field. It can also be used to detect changing strategies of a given trader as well. This is accomplished through a process of outlier detection where correlation based outliers will indicate those transition matrices that are markedly different from the rest of the field.

We also believe that this modeling can give insights to detecting the insider dealing market manipulation as it provides a framework by which to compare between the trading strategies of different traders.

Future work may consider expanding the number of states to capture finer movements in the variables that reflect supply and demand for a security. We also believe that it would be possible to cluster the transition matrices or a reduced dimensional representation of them to find groups of traders whose behaviors are broadly similar. This type of clustering can also be used to detect outliers and thereby anomalous traders and trading strategies.

References

[1].Z. Brezezniak & T. Zastawniak, Basic Stochastic Processes, 1st ed., vol. 1. London: Springer, 2000, pp.85-137.
[2]. F.M. Reza, An introduction to information theory, 1st ed, vol. 1, New York: McGraw-Hill, 1994, pp. 77-130.
[3]. Y. Cao, Y. Li, S. Cleman, A. Belatreche, T..M. McGinnity, Adaptive hidden markov model with anomaly states for price manipulation detection, IEEE Trans. Neural Networks and Learning Systems, vol. 26, pp. 318-330, April 2014.
[4]. F. Allen and D. Gale, "Stock price maniplulation", The Review of Financial Studies, vol. 5, no. 3, pp. 503-529, 1992.
[5]. Matteo Mateucci, A tutorial on Clustering Algorithms: K-Means,"https://home.deib.polimi.it/matteucc/Clustering/tutorial_html/kmeans.html", extracted on March 3, 2017.
[6]. E.M. Miranda, P. McBurney & M. Howard, "Learning unfair trading: a market manipulation analysis from the reinforcement learning perspective", Department of Infomatics, Kings College London.