Best execution compliance automation: towards an equities compliance workstation

Michael Mainelli and Mark Yeandle
Z/Yen Limited, London, UK

Abstract

Purpose – Forthcoming requirements in MiFID and RegNMS mean that buy-side and sell-side firms need to find ways of showing regulators that they are sifting through their trading volumes in a justifiable, methodical manner looking for anomalous trades and investigating them, in order to prove "best execution". The objective was to see if a SVM/DAPR approach could help identify equity trade anomalies for compliance investigation.

Design/methodology/approach – A major stock exchange, a computer systems supplier, four brokers and a statistical firm undertook a cooperative research project to determine whether automated statistical processing of trade and order information could provide a tighter focus on the most likely trades for best execution compliance investigation.

Findings – The support vector machine approach worked on UK equities and has significant potential for other markets such as foreign exchange, fixed income and commodities.

Research limitations/implications – The research has implications for risk professionals as a generic approach to trading anomaly detection. The prototype compliance workstation can be trialed.

Originality/value – Automated anomaly detection could transform the role of compliance and risk in financial institutions.

Keywords United Kingdom, Foreign exchange, Financial institutions, Financial control

Paper type General review

Research project objective

In 2005 a joint research project was agreed between the London Stock Exchange, Sun Microsystems and Z/Yen Limited. The primary objective was to investigate the feasibility of using support vector machine and dynamic anomaly and pattern response (SVM/DAPR) techniques to automate the detection of best execution anomalies for management investigation (Mainelli and Yeandle, 2006). To meet this objective, the research would also compare the results of the SVM/DAPR technique with current techniques such as VWAP and comparisons with the current best price.

The authors would like to thank Alan Helmore-Simpson, the project manager of PropheZy development, and John Murray and Tim Scarfe, the project managers of VizZy development, for all their assistance. They would also like to thank Christopher Hall, Ian Harris and Simon McClean for their quality-assurance advice during the project, as well as Jeremy Horne for tremendous amounts of data processing. In addition, the authors would like tho thank Gervase Clifton-Bligh of Fractal Edge and Freddie McMahon of Decisionality for the use of their software in demonstrating how a compliance workstation might work. They would also like to thank the participants in this research without whose data this study would not have been possible. Finally, they would like to thank the study's sponsors, Sun Microsystems (with a special mention to the ever-supportive Nigel Woodward) and the London Stock Exchange (with a special mention to the ever-helpful Gurbindar Bansal) for their support.

The Journal of Risk Finance
Vol. 7 No. 3, 2006
pp. 313-336

1526-5943
DOI 10.1108/15265940610664988

The research would also evaluate how useful SVM/DAPR techniques are in providing a tighter set of trades for further investigation.

Some of the questions that the research attempted to answer are:

- How large is the universe of anomalous trades?
- Using SVM/DAPR techniques, how many trades actually warrant investigation, and what proportion of the universe do these represent?
- What do firms do now to monitor the execution quality?
- Could a SVM/DAPR approach provide a benchmark for measuring best execution better than VWAP comparisons with best price?

The "null hypothesis" was effectively that automated sifting and selection will be unable to identify potentially anomalous trades any better, if at all, than existing processes.

Approach

Research team

Research team roles were allocated to:

- London Stock Exchange, who provided Execution Quality Service and tick data, as well as help in recruiting members as participants.
- Sun Microsystems, who provided direct financial backing and further services and equipment, particularly Sun Solaris equipment for large volume SVM/DAPR.
- Z/Yen, who provided PropheZy and VizZy, as well as project management and researchers to build the SVM/DAPR systems and conduct the data trials. A Z/Yen director, Michael Mainelli, was the principal point of contact for the research project and Mark Yeandle was the Project Manager.

Overall approach

Following an informal trial in 2004 (Mainelli and Yeandle, 2006), this research project was conducted from June to December 2005 after the team recruited four brokers prepared to examine off-book equity trades. The research began by surveying participants' existing approaches to best execution compliance. The team then built SVM/DAPR systems for each broker taking three months of trading data for training in order to test a fourth month of data for anomalous trades. The research team sent back to each broker a set of trades that the SVM/DAPR system identified as worth being investigated. The key criterion for investigation was that the trade price seemed anomalous. The research team then worked on contrasting the SVM/DAPR output with alternative methods of identifying outliers. These alternatives included:

- VWAP;
- best price at the time of the trade; and
- cluster analysis.

Methodology

The research followed an overall methodology as follows, though many tasks were conducted in parallel:

(1) review of previous relevant research;
(2) recruit four participants;
(3) outline the data requirements;
(4) assess current best execution monitoring practices;
(5) collect data from the brokers and the LSE;
(6) analyze, validate, and pre-process the data;
(7) build trial SVM/DAPR systems for various parameters to ensure historic data has predictive capacity;
(8) select approach to anomaly identification and build full SVM/DAPR systems;
(9) develop visualizations of a possible compliance workstation;
(10) feed back data and visualizations to the participants;
(11) collect feedback from the participants; and
(12) prepare final research document for publication.

Participants
The four brokers who participated in the research provided:

- Staff time to explain current best execution compliance approaches, factors, coefficients and benchmarks.
- A dataset of four months' off-book and on-book equities trading data (September to December 2004).
- Staff time to investigate a sample of trades that the SVM/DAPR approach found anomalous.

The research team agreed to keep the identities of the participants confidential. Their profiles are:

- Broker A is an independent UK-based firm offering stockbroking services, asset management and financial planning advice to members of the public.
- Broker B is a US-based global investment bank serving corporations, institutions, governments and high-net-worth investors worldwide.
- Broker C is an Asian-based financial services group engaged in four main business areas –Global Markets, Investment Banking, Merchant Banking and Asset Management.
- Broker D is a European-based global investment bank and asset management business dealing with individuals, corporations, the public sector and other not-for-profit organisations.

In return for participating in this research the brokers received:

- Research results including working documents not subject to confidentiality.
- A presentation of the approach and results for their business.
- An example of a SVM/DAPR system on their data and sample visualization of their data in a prototype compliance workstation.

PropheZy, a support vector machine/dynamic anomaly and pattern response (SVM/DAPR) implementation

This study used classification and prediction tools based on SVM mathematics to undertake predictive analysis of the data. SVMs are algorithms that develop classification and regression rules from data. SVMs result from classification algorithms first proposed by Vladimir Vapnik in the 1960s, arising from his work in Statistical Learning Theory (Vapnik, 1995, 1998). SVMs are based on some wonderfully direct mathematical ideas about data classification and provide a clear direction for machine learning implementations. While some of the ideas behind SVMs date back to the 1960s, computer implementations of SVMs did not arise until the 1990s with the introduction of a computer-based approach at COLT-92 (Boser *et al.*, 1992).

SVMs are now used as core components in many applications where computers classify instances of data (e.g. to which defined set does this group of variables belong), perform regression estimation and identify anomalies (novelty detection). SVMs have been successfully applied in time series analysis, reconstructing chaotic systems, and principal component analysis. SVM applications are diverse, including credit scoring (good or bad credit), disease classification, handwriting recognition, image classification, bioinformatics, and database marketing, to name a few.

SVMs are said to be independent of the dimensionality of feature space as the main idea behind their classification technique is to separate the classes in many data dimensions with surfaces (hyperplanes) that maximize the margins between them, applying the structural risk minimization principle. The data points needed to describe the classification algorithmically are primarily those closest to the hyperplane boundaries, the "support vectors." Thus, only a small number of points are required in many complex feature spaces. SVMs can work well with small datasets, though the structure of the training and test data is an important determinant of the effectiveness of the SVM in any specific application.

SVMs compete forcefully with neural networks as well as other machine learning and data mining algorithms as tools for solving pattern recognition problems. Where SVMs do not perform well it is arguable that the algorithmic rules behind the support vector algorithm do not so much reflect inabilities of the learning machine (as in the case of an over fitted artificial neural network) as much as irregularities of the data. In short, current opinion holds that if the data in the domain is predictive, SVMs are highly likely to be capable of producing a predictive algorithm. Importantly, SVMs are robust tools (understandable implementations, simple algorithmic validation, better classification rates, over fitting avoidance, fewer false positives and faster performance) in practical applications. "The SVM does not fall into the class of "just another algorithm" as it is based on firm statistical and mathematical foundations concerning generalization and optimization theory" (Burbridge and Buxton, 2001). However, comparative tests with other techniques indicate that while SVMs are highly likely to be capable of predicting, in some applications SVMs may not be the best approach for any specific dataset. "In short, our results confirm the potential of SVMs to yield good results, especially for classification, but their overall superiority cannot be attested" (Meyer *et al.*, 2002).

PropheZy and VizZy are two software packages developed by Z/Yen for classification and visualization of data (www.zyen.com/Products/Prophezy/

prophezyhtm, www.zyen.com/Products/Vizzy/vizzy.htm). Together they constitute a SVM/DAPR environment. PropheZy implements a SVM on a server (though it can be used in a local client/server mode). Naturally, as in any field of computing, there are a number of variant SVM implementations, of which PropheZy implements three types – C-SVC, nu-SVC, and binary. Further, of statistical importance in replicating results is the "kernel function." PropheZy implements four types of kernel function – linear, radial basis function, sigmoid, and polynomial. The SVM types and kernel function types are described in detail in Vapnik (1995, 1998). For this study, the SVM implementation used was C-SVC and the kernel function was linear.

So far, the PropheZy server SVM has been implemented on a Linux server, a Sun Solaris server and a Windows NT server. PropheZy implements the user-interface to the server SVM via XML (extensible mark-up language). The XML user-interface can be via an HTML page, directly through a bulk file loader command line or by use of an Excel add-in that performs XML data submission from spreadsheets to the server SVM and displays results back in Excel spreadsheets. For this study, the PropheZy implementation was the Sun Solaris server using a bulk file loader command line. VizZy provided visualization of clustering, histogram, Voronoi, and data cube diagrams from tabular data output by PropheZy.

Z/Yen has benchmarked PropheZy against standardized machine learning tests, e.g. appropriate StatLog test sets (Michie *et al.*, 1994), in order to validate the SVM with good to excellent results. Z/Yen has trialed PropheZy extensively in financial services applications and sees great promise for SVMs and other Dynamic Anomaly and Pattern Response (DAPR) techniques in areas such as compliance (Mainelli, 2005), trade anomaly detection and scorecards (Mainelli, 2004) as well as regression and value prediction (Mainelli *et al.*, 2003).

Data

For the purposes of this research, a discrete, easily quantifiable, and readily obtainable set of data was required. SETS – The London Stock Exchange's trading service for U.K. blue chip securities is an electronic order book that can execute hundreds of trades a second. Securities traded on SETS include all the FTSE 100 securities, the most liquid FTSE 250 securities along with some others. The prices quoted on SETS are the best available at the time and as such trades fulfilled "on-book" might be viewed as representing best execution. However, SETS-traded equities are also regularly traded outside the electronic order book ("off-book" trades).

The research team decided that they should look at off-book trades of SETS securities, as this was the area requiring proof of best execution. A large number of trades are conducted off-book and it is often these trades that attract attention from compliance departments. There are a substantial number of these trades that occur outside the current bid/offer spread. This is usually for a very good reason – typically the size of the trade (in relation to the usual size of trades in that security) or specific client instructions regarding timing of execution.

The total number of trades included within the research was approximately 190,000 with a value of over £54 billion. These were from four brokers and covered their off-book trades from September to December 2004 inclusive. It should be noted that although exchange-traded securities were selected as an ideal dataset for this research, the information used in building the model was only that which would be available for non-exchange-traded instruments.

Data preparation and validation

Each of the four participants provided large text files containing details of their off-book trades for the four months from September to December 2004 inclusive. The first thing that the research team did was to validate these data with the data provided by LSE. The team also studied the movement of the market in general over the period (Figure 1) in order to establish overall volume and price patterns.

From the data supplied by all parties, the team generated files of transactions with the following fields (* indicates that the data was codified into numeric values from text):

- SEDOL code (share code, e.g. 0316893 = Eurotunnel)*;
- market segment of share (e.g. pharmaceutical, retail)*;
- trade date;
- trade time;
- trade size;
- trade price;
- trade code;
- buy/sell indicator (buy = 1, sell = 0)*;
- participant code buyer (counterparty)*;
- participant code seller (counterparty)*; and
- settlement due date.

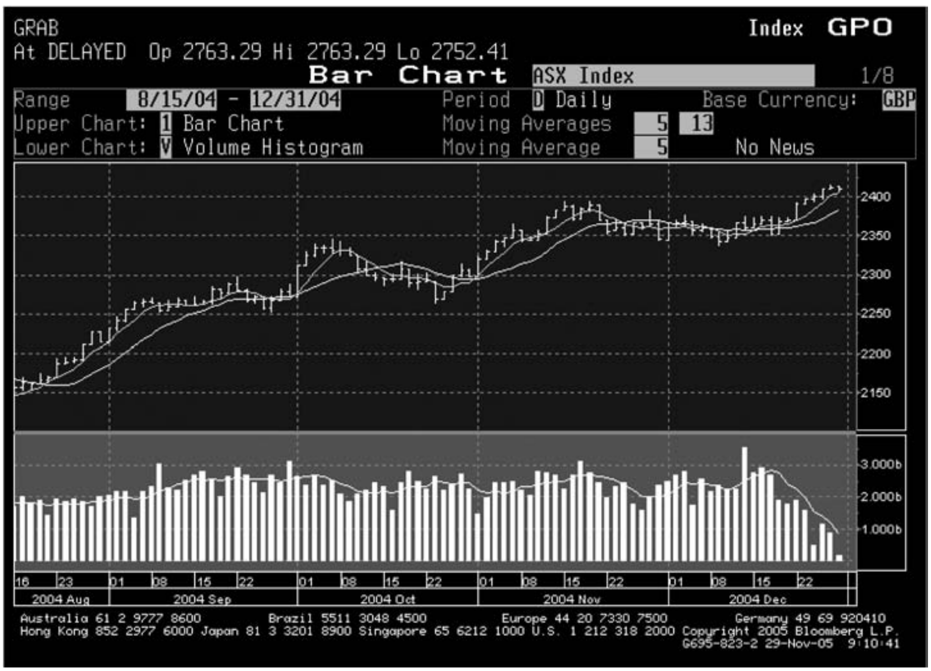

Figure 1.
Movement of the all share index, September to December 2004

The following fields were then added:

- market sector (e.g. FTSE 100, FTSE 250 – downloaded from LSE web site)*;
- day of the week (1 to 5 – derived from date);
- closing price for the previous ten days;
- inside (0) or outside (1) the bid/offer spread;
- consideration (a calculated field – price multiplied by the quantity of shares);
- five-day percentage price movement (calculated fields);
- bid price (from LSE files – linked using SEDOL Code);
- mid price (from LSE files – linked using SEDOL Code);
- offer price (from LSE files – linked using SEDOL Code);
- VWAP movement (calculated previous trades executed by each firm within the previous five days);
- FTSE Index movement (of the relevant index for the share) since the last trade;
- price volatility –the standard deviation of the closing prices of the previous 12 days (12 days based on discussions with market participants and previous research);
- Percentage size of bid/offer spread;
- liquidity (average number of shares traded per day, over the 88 trading days between September and December);
- return versus mid price (trade price – mid price)/mid price;
- return versus previous closing price (trade price – previous closing price)/previous closing price;
- percentage of liquidity of the trade (trade size/liquidity);
- three-day index movement for the index in which the share is listed;
- number of trades in a day, for each share and each day;
- time since last trade – the number of ten minute slots since the previous trade of that share by that broker; and
- day of the week of the last trade (one to five) – when the share was last traded by that broker.

Compliance workstation prototype

During the course of the project the team built a prototype "Compliance Workstation." The Compliance Workstation combined a number of software tools, PropheZy, VizZy, FractalEdge, and Decisionality within an Excel framework. The Compliance Workstation provided a number of features, specifically the ability to:

- Construct predictive tests on any trade characteristic in order to spot anomalies.
- Spot anomalies using cluster analysis.
- Display the results visually, specifically showing predicted versus actual differences in three dimensions (Figure 2 shows a visualisation with the predicted price movement bands in dark dots plotted against the actual price movement band in lighter dots. The length of the white link shows the difference between the actual and predicted values).

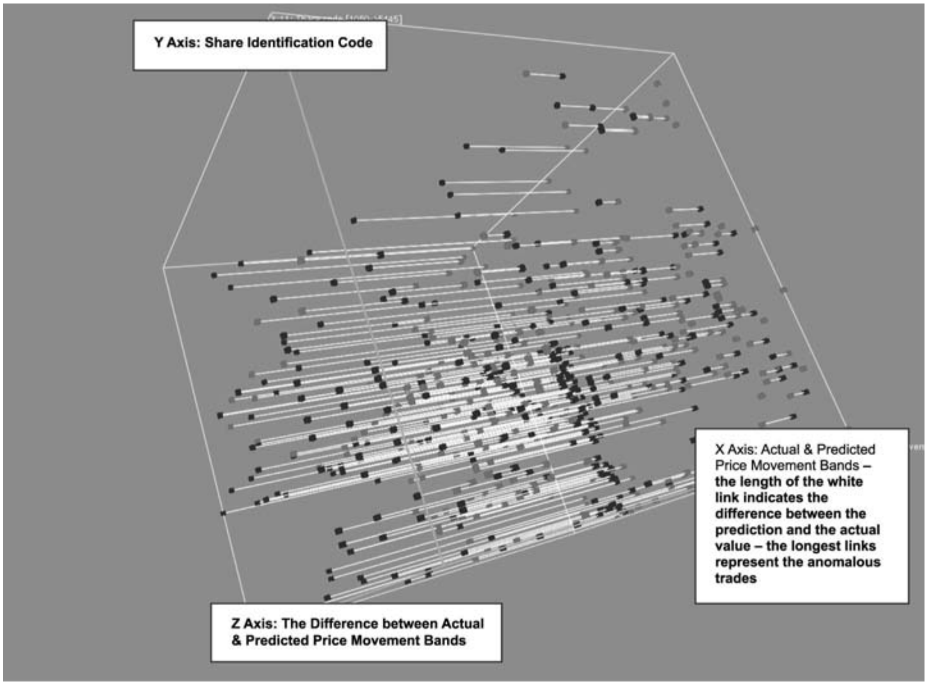

Figure 2.
A three-dimensional VizZy presentation of anomalous trades

- Provide a "drill down" tool for a compliance officer to home in on specific trades. This tool, Fractal Intelligence, allows the user to drill down through a variety of hierarchies. Figure 3 shows a set of sell trades that fell outside the bid/offer spread and is arranged in ten circles, starting clockwise from the top showing trades with increasing differences between actual and predicted price movement bands.

Initial tests of predictive capability

The research team initially conducted a range of tests on the data to assess the predictive capacity of the data. Many of the heuristics used in the construction of these tests were informed by previous studies of share price prediction (Cao and Tay, 2001, 2003 and Huang *et al.* 2005), as well as previous Z/Yen client work on share liquidity and other analyses.

The first set of tests used the SVM to predict the counterparty with which each transaction was conducted. This was conducted on two datasets:

- In the first dataset there were 52 different counterparties and the SVM was able to correctly predict the counterparty in over 43 percent of the trades tested.
- In the second dataset there were 11 different counterparties and over 61 per cent were correctly predicted.
- The second set of tests used the SVM to predict the share that was traded in each transaction. This was again conducted on two datasets: In the first dataset there

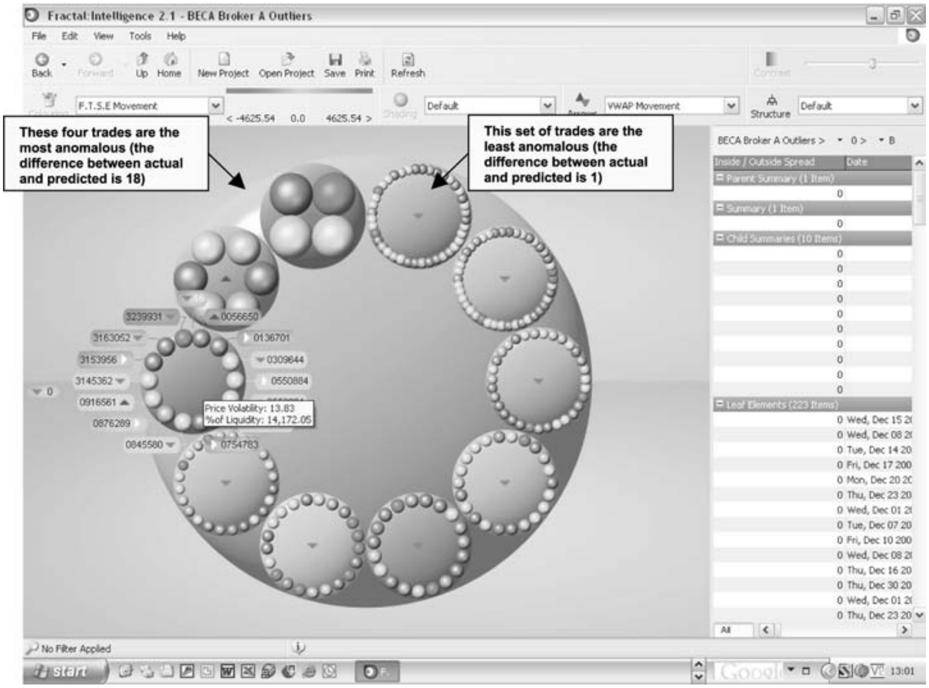

Figure 3.
A "drill-down" and data visualization tool showing anomalous trades

were 150 different shares and the SVM was able to correctly predict the share in nearly 8 percent of the trades tested.

- In the second dataset there were 82 different shares and over 18 percent were predicted correctly.

One question was whether the predictions made by the SVM were different to those that could have been made by a simpler method, a "naïve classifier." The team compared the SVM output with a random classification based on the observed probabilities of each class. This was done by conducting a statistical significance test based on Liddell's exact test for paired proportions (Liddell, 1983), which examines the difference between the two alternatives by looking at where each method was correct when the other was incorrect. The tests show that for Broker's A, B and D, there is significant evidence that the predictions made by the SVM are much better than those that would be made by a naïve classifier. The evidence for Broker C (a much smaller dataset) was weaker, but nevertheless supported the performance of the SVM.

A third set of tests run, were to see whether the SVM could predict how much the VWAP of a share would move. For this set of tests, VWAP movement was banded into twenty equal bands across the range of VWAP, between −0.3 per cent and +0.3 per cent movement, with each band having a size of 0.03 per cent. Tests conducted by building a SVM with three months' data were significantly more accurate than those built with only one month's data. Increasing the dataset to four months gave no significant increase in accuracy (5-8). The team therefore ran all datasets using a three-month history to train the SVM. The average accuracy over all four brokers was over 25 per cent. Over 52 per cent of predictions were within four bands.

The conclusion was that the data did have predictive capacity when used by a SVM.

Results and analysis

Price movement band tests

After conducting a number of trial tests and speaking to several industry experts, the research team decided that the most relevant test was to predict price movement. In order to keep the SVM construction and the analysis simpler, the team calculated 20 price movement bands. First, the team calculated the percentage by which a share price had moved between the trade in question and the previous trade. Then the team calculated the natural logarithm of the absolute price movement. Using the mean and standard deviation of all price movement logarithms in the training set, the team created a normalized price movement variable by subtracting the mean value from each observation and dividing by the standard deviation. Based on these normalized variables, the trades were split into 20 equal price movement bands, with the probability of any trade being in any particular band at 5 per cent. By using this method of banding, the trades with the largest share price movements fall into the highest bands and the trades with the smallest movements are contained in the lowest bands.

Tests were then conducted using three months of data to train the model and perform price movement band predictions for trades on the following day (the first test used September 1 to November 1 as a training set and predicted December 1; the second test used September 2 to December 1 as a training set and predicted December 2, etc.). This "rolling" approach covered the 21 trading days of December 2004.

The team looked at the predictions of price movement band and compared these to the actual price movement band. The overall results were that over 9 per cent of records were predicted correctly, 47 per cent were within four bands, 74 per cent were within nine bands and 93 percent were within 14 bands. After examination and discussion with industry experts, the trades where the prediction differed from the actual result by 15 bands or more were defined as "outliers". On average, 7 per cent of trades were statistically anomalous or "outliers". If all statistically anomalous trades that were transacted at "best" price (or better) are excluded, only 1 per cent of all trades remain for possible investigation. The daily results for each of the four brokers are shown in both tabular (Tables I-IV) and graphical form (Figures 3-6). The accuracy of the SVM at correctly predicting the exact price movement band out of 20 varies between 8 per cent and 11.75 per cent for the different brokers. The accuracy of the SVM at predicting within four price movement bands varies between 45 and 56 per cent for the different brokers (Figure 7).

The team was keen to establish that these results were better than those of a "naïve classifier" and conducted Liddell's test for paired proportions. Again, the predictions made for Broker's A, B, and D are significantly better than those that would be made by a naïve classifier. The evidence for Broker C (a much smaller dataset) was weaker, but still supported the performance of the SVM. The team also put in a handful of "wacky" trades, i.e. trades where the variables were copied from real trades, but the

Table I. Analysis of the 21 price movement band tests for broker A (21 trading days in December 2004)

		Broker A Kits							
Testdate	Correct	Accuracy (%)	Within 0-4 (%)	Within 0-9 (%)	Within 0-14 (%)	Outliers (%)	MSE (%)	R-squared (%)	Number of predicted bands
01/11/2004	10 of 110	9.09	51.82	84.55	95.45	5	45.25	25.00	10
02/12/2004	10 of 104	10.57	51.92	76.92	96.15	4	52.08	16.00	8
03/12/3004	13 of 109	11.93	56.88	78.90	91.74	9	51.40	16.00	9
06/12/2004	9 of 113	7.96	55.75	81.42	97.35	3	43.79	29.00	10
07/12/2004	14 of 108	12.96	48.15	72.22	97.22	3	55.46	11.00	10
08/12/2004	17 of 160	10.63	51.88	78.75	95.63	7	51.98	17.00	13
09/12/2004	2 of 45	4.44	55.56	77.78	95.56	2	52.00	18.00	9
10/12/2004	20 of 125	16.00	55.20	79.20	92.00	10	51.07	24.00	10
13/12/2004	12 of 97	12.37	49.48	80.41	95.85	4	49.11	18.00	14
14/12/2004	20 of 145	13.79	59.31	73.10	94.48	8	53.55	21.00	12
15/12/2004	15 of 141	10.64	61.70	82.27	95.74	6	41.89	33.00	II
16/12/2004	27 of 145	18.62%	56.55	80.69	96.55	5	43.77	36.00	10
17/12/2004	14 of 124	11.29	57.26	82.26	94.35	7	47.98	34.00	10
20/12/2004	19 of 206	9.22	50.97	79.13	94.66	11	50.51	27.00	10
21/12/2004	14 of 154	9.09	56.49	82.47	94.81	8	45.25	33.00	10
22/12/2004	14 of 114	12.28	52.63	79.82	97.37	3	50.18	26.00	8
23/12/2004	2 of 74	2.71	44.59	79.73	95.95	3	49.24	34.00	9
24/12/2004	6 of 34	17.64	58.82	79.41	94.12	2	46.18	32.00	8
29/12/2004	4 of 51	7.84	41.18	78.43	92.16	4	53.84	16.00	10
30/12/2004	8of 50	16.00	56.00	76.00	96.00	2	47.56	36.00	7
31/12/2004	5 of 23	21.74	56.52	78.26	86.96	3	58.30	18.00	7
Average		11.75	53.75	79.13	94.77	5.19	49.54	24.29	9.76

Table II.
Analysis of the 21 rice movement band tests for broker A (21 trading days in December 2004)

					Broker B tests				
Tesi Date	Correct	Accuracy (%)	Within 0-4 (%)	Within 0-9 (%)	Within 0-14 (%)	Outliers (%)	MSE (%)	R-squared (%)	Number of predicted bands
01/12/2004	27 of 324	8.33	51.23	75.31	94.75	17	52.58	7.06	13
02/12/2004	63 of 595	10.59	59.16	B2.18	95.46	27	44.73	12.66	16
03/12/2004	37 of 458	8.08	52.62	82.97	96.72	15	43.18	9.26	14
06/12/004	30 of 333	9.01	60.06	87.69	96.10	13	37.28	21.34	16
07/12/2004	32 of 396	8.08	45.48	78.28	94.95	20	50.43	8.68	16
08/12/2004	22 of 311	7.07	50.16	76.21	93.89	19	54.72	8.91	12
09/12/2004	15 of 306	4.90	46.08	74.18	94.12	18	58.91	3.09	13
10/12/2004	36 of 469	7.68	56.72	81.88	97.44	12	42.99	23.83	16
13/12/2004	45 of 388	11.60	57.47	85.05	95.62	17	41.57	18.45	14
14/12/2004	23 of 386	5.96	46.63	72.28	90.41	37	64.48	2.25	12
15/12/2004	31 of 413	7.51	51.57	77.00	94.92	21	52.30	11.49	12
16/12/2004	21 of347	6.05	50.72	75.79	94.81	18	53.97	5.25	16
17/12/2004	20 of 337	5.93	50.15	77.45	93.47	22	55.40	5.97	14
20/12/2004	20 of 223	3.97	57.40	78.48	96.86	7	46.66	11.95	14
21/12/2004	31 of 351	8.83	49.00	76.92	95.16	17	53.58	7.17	13
22/12/2004	24 of 233	10.30	49.79	79.40	96.14	9	50.23	6.77	15
23/12/2004	18 of 236	7.63	52.97	79.66	95.76	10	51.04	8.22	13
24/12/2004	2 of 30	6.67	63.33	93.33	100.00	10	28.07	8.58	8
29/12/2004	21 of 162	12.96	61.11	82.10	95.06	8	43.53	21.64	13
30/12/2004	7of 154	4.55	48.70	87.01	98.05	3	41.66	10.65	13
31/12/2004	6 of 78	7.69	65.38	92.31	97.44	2	30.18	38.20	10
Average		8.02	53.75	80.74	95.58	14.86	47.50	11.97	13.48

actual price movements were artificially exaggerated or muted. The SVM correctly identified four out of the five as anomalous.

Filtering outliers

An outlier, or anomalous trade, was defined as a trade where the predicted price movement differs from the actual price movement by more than 15 bands out of 20 – either a very high price movement was predicted but a low price movement was observed, or a very low price movement was predicted but a high price movement was observed.

Table V indicates that when using the SVM as a filter (First Filter) on average 7 per cent of non-SETS trades are defined as outliers. There are still too many outliers for a detailed manual investigation. A second filter is therefore needed. A trade is unlikely to fail best execution if it was conducted at the best prevailing price (or better), though there are some arguments that very large trades might be capable of exceptional improvement under certain conditions. When excluding trades outside the bid/offer spread (Second Filter) is combined with the first filter the number of outliers that are outside the bid/offer spread is approximately 1 per cent.

To verify further the results shown in Tables I-IV, the team analyzed the quality of predictions within the "extreme" bands (Table VI). These are the three bands with the lowest price movement (0, 1 and 2) and the three with the highest price movements (17, 18 and 19).

Table III. Analysis of the 21 price movement band tests for broker A (21 trading days in December 2004)

Test date	Broker C tests Correct	Accuracy (%)	Within 0-4 (%)	Within 0-9 (%)	Within 0-14 (%)	Outliers (%)	MSE (%)	R-squared (%)	Number of predicted bands
01/12/2004	0of 14	0.00	21.43	71.43	100.00	0	64.50	6.83	8
02/12/2004	2of 15	13.33	60.00	86.67	100.00	0	31.93	28.15	5
03/12/2004	1of 13	7.69	53.85	92.31	100,00	0	30.31	3.44	4
06/12/2004	3of 13	23.08	76.92	92.31	100.00	0	18.85	30.01	5
07/12/2004	1 of 15	6.67	60.00	80.00	100.00	0	40.93	11.20	4
08/12/2004	3 of 23	13.04	56.52	86.96	95.65	1	38.04	11.86	*6*
09/12/2004	1 of 10	10.00	30.00	90.00	90.00	1	47.80	15.00	3
10/12/2004	8 of 27	29.63	59.26	92.59	96.30	1	36.04	13.58	8
13/12/2004	0 of 14	28.57	50.00	85.71	92.86	1	49.00	14.43	5
14/12/2004	1 of 16	6.25	93.75	100.00	100.00	0	8.56	90.53	4
15/12/2004	3 of 19	15.79	57.89	89.47	94.74	1	38.95	4.44	6
16/12/2004	2 of 20	10.00	65.00	90.00	100.00	0	34.55	10.53	4
17/12/2004	2of 16	12.50	87.50	100.00	100.00	0	9.38	35.31	4
20/12/2004	1 of 19	5.26	47.37	84.21	100.00	0	53.16	0.01	4
21/12/2004	2 of 15	13.33	66.67	86.67	93.33	1	39.00	8.47	4
22/12/2004	1 of 12	8.33	41.67	83.33	100.00	0	41.67	5.76	4
23/12/2004	0of 11	0.00	45.45	100.00	100.00	0	16.00	20.70	5
24/12/2004	1 of 5	20.00	60.00	80.00	100.00	0	60.40	41.54	2
29/12/2004	0of 3	0.00	66.67	66.67	100.00	0	62.00	21.88	3
30/12/2004	0of 6	0.00	33.33	10000	100.00	0	31.17	7.60	2
31/12/2004	1 of 8	12.50	62.50	100.00	100.00	0	14.50	26.94	4
Average		11.14	56.94	88.49	98.23	0.29	36.51	19.44	4.48

Table IV.
Analysis of the 21 price movement band tests for broker A (21 trading days in December 2004)

Test date	Correct	Accuracy (%)	Within 0-1 (%)	Within 0-9 (%)	Broker D tests Within 0-14 (%)	Ouiliers (%)	MSE (%)	R-squared (%)	Number of predicted bands
01/12/2004	126 of 1477	8.53	43.60	68.99	91.47	126	68.44	9.67	8
02/12/2004	129 of 1471	8.77	42.15	69.54	90.55	139	71.33	7.61	9
03/12/2004	118 of 1355	8.71	45.46	70.26	91.51	115	67.37	10.83	8
06/12/2004	137 of 1300	10.54	48.85	74.08	92.38	99	60.52	14.82	10
07/12/2004	165 of 1809	9.12	42.18	73.08	93.26	122	62.95	12.71	10
08/12/2004	168 of 1748	9.61	43.99	70.48	92.68	128	66.41	12.71	9
09/12/2004	142 of 1605	8.85	41.43	69.84	91.84	131	68.14	8.73	12
10/12/2004	115 of 1287	8.94	43.90	71.56	91.22	113	65.69	10.43	10
13/12/2004	151 of 1766	8.55	43.77	70.27	91.51	150	66.81	9.73	12
14/12/2004	163 of 1496	10.90	43.92	70.72	91.58	126	66.29	11.18	10
15/12/2004	104 of 1345	7.73	42.16	71.08	91.53	114	67.50	9.68	10
16/12/2004	135 of 1624	8.31	47.97	73.03	93.10	112	61.95	16.24	9
17/12/2004	142 of 1661	8.55	45.33	70.20	90.19	163	69.31	8.70	9
20/12/2004	121 of 1348	8.98	44.36	73.74	92.21	105	63.16	14.28	10
21/12/2004	129 of 1473	8.76	41.00	68.70	91.85	120	69.30	8.71	11
22/12/2004	193 of 1901	10.15	46.24	72.33	92.79	137	63.31	13.80	10
23/12/2004	122 of 1251	9.75	46.52	72.42	95.12	61	58.47	17.38	9
24/12/2004	49 of 480	10.21	50.83	76.67	95.83	20	52.47	24.49	9
29/12/2004	119 Df935	12.73	52.94	73.37	91.76	77	59.66	16.29	8
30/12/2004	115 of 845	13.61	48.17	74.08	94.12	47	56.90	19.44	9
31/12/2004	42 of 446	9.42	54.71	82.06	96.64	15	46.03	23.48	9
Average		9.56	45.69	72.21	91.58	105.71	63.43	13.38	9.67

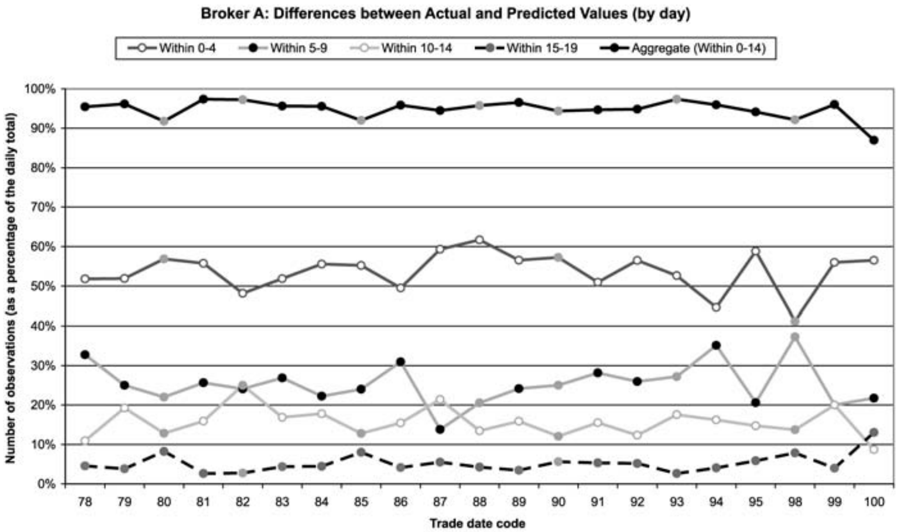

Note: The dashed line shows the percentage of "Outliers", trades where the difference is 15 or more

Figure 4.
Differences for broker A between actual and predicted price movement bands – out of 20 Bands (yellow line shows the percentage of "outliers", trades where the difference is 15 or more)

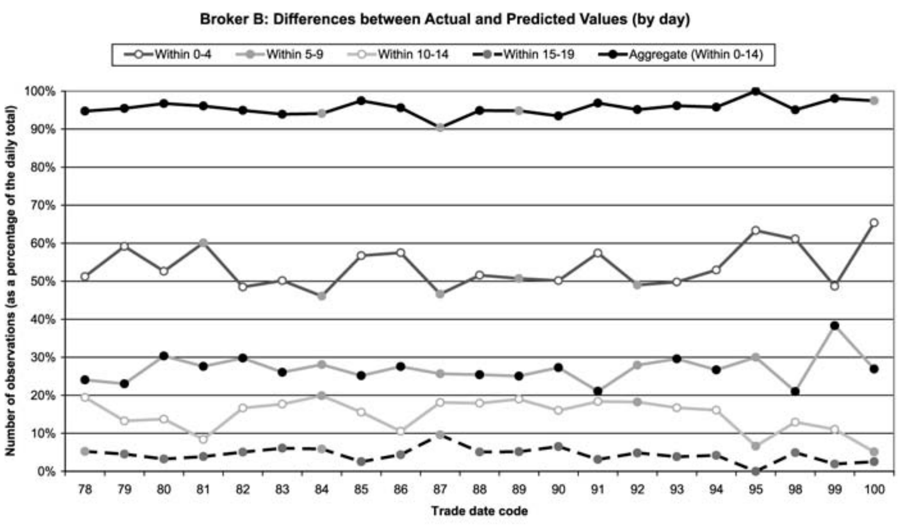

Note: The dashed line shows the percentage of "Outliers", trades where the difference is 15 or more

Figure 5.
Differences for broker B between actual and predicted price movement bands – out of 20 bands (yellow line shows the percentage of "outliers", trades where the difference is 15 or more)

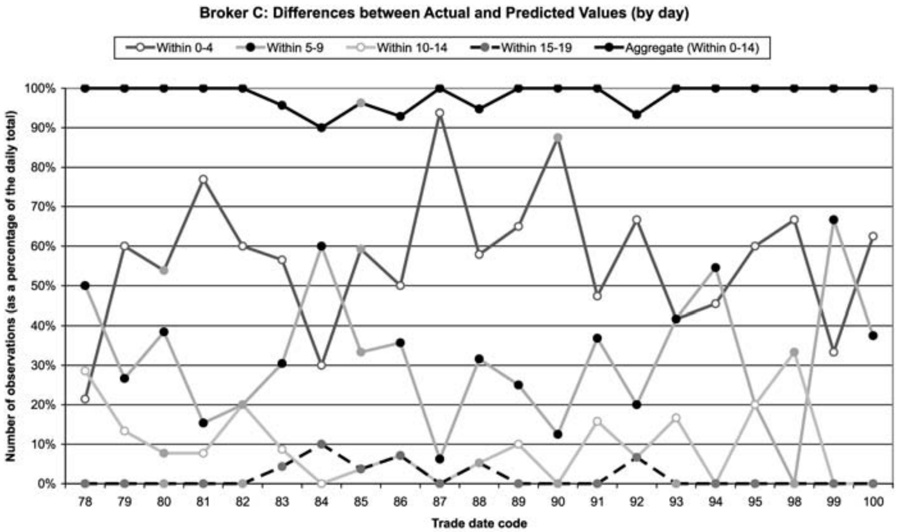

Note: The dashed line shows the percentage of "Outliers", trades where the difference is 15 or more

Figure 6.
Differences for broker C between actual and predicted price movement bands – out of 20 bands (yellow line shows the percentage of "outliers", trades where the difference is 15 or more)

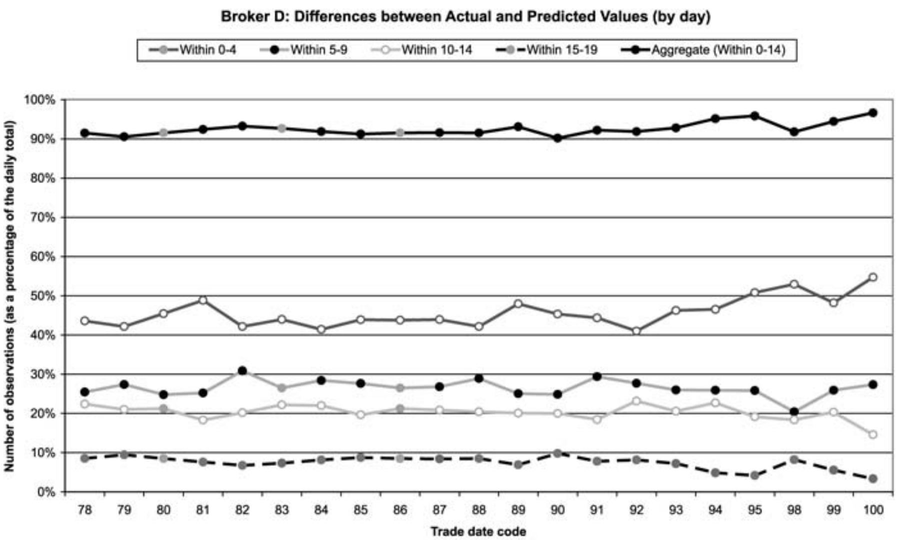

Note: The dashed line shows the percentage of "Outliers", trades where the difference is 15 or more

Figure 7.
Differences for broker D between actual and predicted price movement bands – out of 20 bands (yellow line shows the percentage of "outliers", trades where the difference is 15 or more)

Table V. An analysis of outliers by broker

Broker	Number of December trades	First filter: Number of outliers	First filter: Percent outliers	Second filter: Number of trades outside bid/offer	Second filter: Percent outside bid/offer	Combined filters: Outliers outside bid/offer	Combined filters: Percent outliers outside bid/offer
A	2,232	109	4.88	56	2.51	1	0.04
B	6,530	312	4.78	2,879	44.09	124	1.90
C	294	6	2.04	11	3.74	1	0.34
D	28,623	2,220	7.76	2,621	9.16	277	0.97
Overall	37,679	2,647	7.03	5,567	14.77	403	1.07

It can be seen from this table that 22.8 per cent of predictions within the six extreme bands were correct and 70.8 per cent were within four bands. Only 16.6 per cent of predictions from these bands were defined as outliers.

There was no outstanding single feature about the trades that were identified as outliers. The team wanted to assess how the outliers identified by the SVM method differed from outliers identified by the two most common methods in use at present, comparison with VWAP and comparison with current market price. An inspection of the list of outliers showed very little crossover with either of these methods.

The average VWAP movement of the 2,647 SVM outliers is 0.28 per cent. The average VWAP movement of the 2,647 trades with the highest VWAP movement is 4.68 per cent – nearly 17 times higher. In a random nine-day sample of trades, the SVM outliers are plotted against ten bands of increasing VWAP movement (Table VII), there is no significant correlation:

The average distance from best market price at the time of trade (from the bid price for sells and the offer price for buys) of the 2,647 SVM outliers is 5.43 per cent. The average distance from best market price of the 2,647 trades with the highest distance is 18.75 per cent – nearly 3.5 times higher. If the SVM outliers are plotted against distance from best market price, there is again no significant correlation (Table VIII).

The team also briefly compared the SVM outliers with a set of outliers produced by single-link hierarchical clustering. A random sample of trades was clustered and the outliers identified. The result of this clustering is shown in Figure 8. The most similar

Table VI.
An analysis of outliers from the extreme bands – three lowest and three highest bands

	Actual band	Correct	Within 0-4	Within 5-9	Within 10-14	Outliers	Total
Low price movements	0	1,034	1,143	84	84	712	2,023
	1	59	867	74	202	415	1,558
	2	18	1,041	72	612	83	1,808
High price movements	17	448	1,510	28	56	247	1,841
	18	226	1,343	30	26	181	1,580
	19	588	1,473	34	8	88	1,603
	Overall	2,373	7,377	322	988	1,726	10,413
	Percent	22.8	70.8	3.1	9.5	16.6	100

Table VII.
An analysis of SVM outliers by band of increasing VWAP

Bands of increasing VWAP	1	2	3	4	5	6	7	8	9	10
Percentage of SVM outliers	13	6	8	6	7	10	13	10	23	4

Table VIII.
An analysis of SVM outliers by band of increasing distance from best price

Bands of increasing distance from best price	1	2	3	4	5	6	7	8	9	10
Percentage of SVM outliers	8	9	11	8	17	12	11	8	9	6

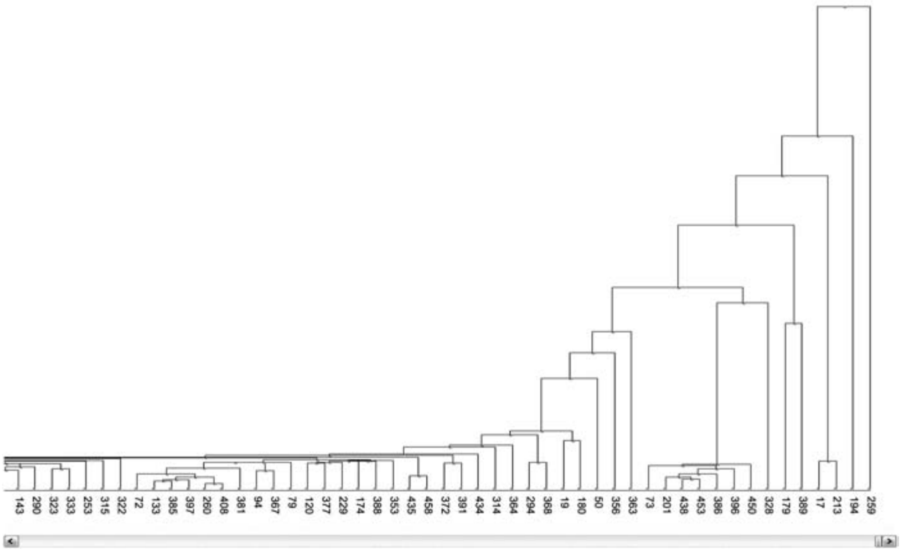

Figure 8.
Clustering of trades (with the most similar trades linked at the bottom and the "outliers" at the far right)

trades, are linked by a horizontal line, towards the bottom of the vertical axis (and left on the horizontal axis). Trades linked higher up (and further right) are less similar to one another. None of the outliers corresponded with the SVM outliers from the same data set; however, the team intends to conduct further research in this area.

One question that was raised during the research was how well did the SVM perform for trades in a single share. Several shares with a reasonable number of trades were examined. The overall level of outlier prediction and distribution of differences between actual and predicted price movement band is very similar to datasets over all shares (Figure 9).

Figure 9 demonstrates that during a time of rapid change in market prices it appears that the SVM does seem to note turning points between trend and uncertain, and uncertain and trend.

The team also examined how well the SVM could predict the magnitude of a single share's price movement (Figure 10). By plotting the actual share price movement since the last trade against the SVM predicted price movement band it is evident that the SVM does achieve reasonable predictions of price movement and ultimately responds well to changes in the market. The three outliers for this frequently traded share are plotted in green. It is interesting to note that in all three cases the SVM presages major changes, but for this broker they had not occurred at that point. The team would recommend further research in this area.

Another aspect of the results that interested the team was the rate of decay of accuracy over time, as it would determine how often a predictive model would need to be rebuilt for any installation of an automated compliance system. The majority of the tests conducted were making predictions for the next full day. It is noticeable that predictions for the afternoon are less accurate than those for the morning. Tests were conducted using half-day model rebuilds (as opposed to whole day). Accuracy did improve marginally using half-day tests. Then tests were conducted using one and two-hour time slots. Hourly and two-hourly tests proved not to be significantly more accurate than the half-day tests. It would appear from the limited amount of testing done in this area that a half-day rebuild would be sufficient, but the team recommends further testing on rebuild intervals.

Usefulness of results

Participant assessment

The team discussed the outliers or anomalous trades with the brokers in order to assess their reaction to the trades that were identified. Some of the remarks made were:

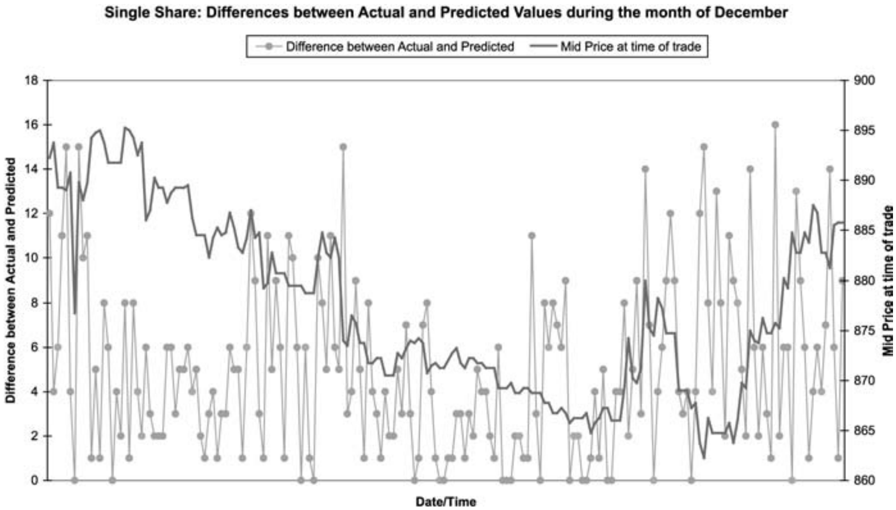

Figure 9.
Difference of actual and predicted price movement band movement for a sngle share plotted against current share price

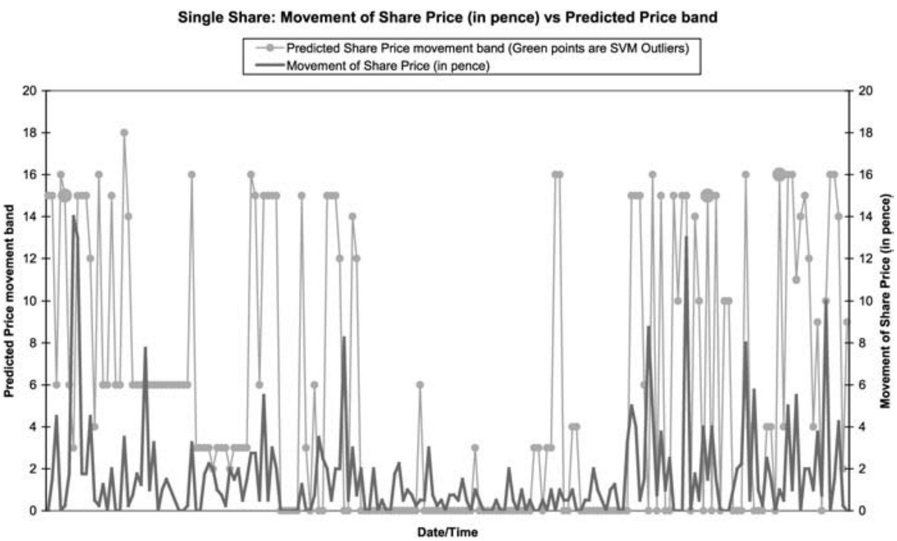

Figure 10.
Movement of share price (in pence) and predicted price movement band for a single share

This system highlighted anomalous trades and if we had been using it, we would certainly have investigated any of these trades that were outside the Bid/Offer spread.

This system would be a great way of seeing a small number of "odd-looking" trades that we could check – the fact that the same principles could be applied to fixed income and other instruments makes it particularly interesting.

An extremely interesting and innovative approach – the trades that were highlighted would certainly have been investigated. Whilst there was no real pattern or similarity in the outliers they were all things that should have been looked at.

Of the trades that were highlighted as anomalous and outside the spread, a compliance officer said, "one was a proprietary trade, but it is very interesting that it was picked up."

Another was part of a convertible bond swap, and the compliance officer said: "this one would definitely have warranted investigation at the time."

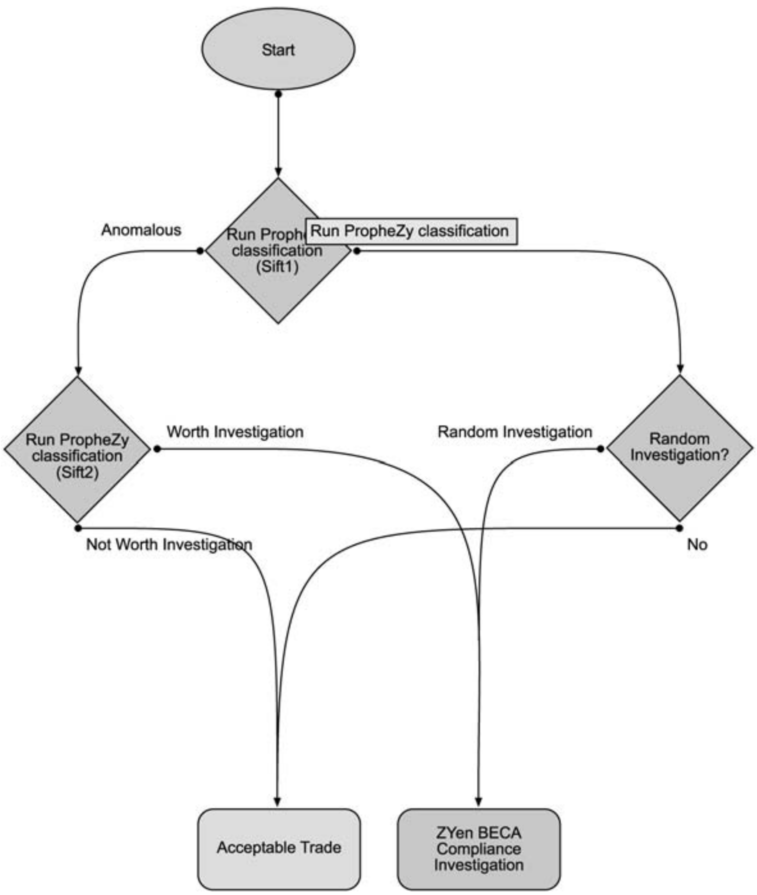

Figure 11.
Possible workflow management tool. Part 1 – identifying trades for investigation

I was fascinated to see the selection of trades that this system identified – there were good reasons why all of them traded at the prices shown but they were just the sort of trades that we should have been looking at.

Several of the trades you picked up on were "in-house" trades – typically closing out short-term positions. Although they were interesting because of the price, we would probably want to filter out most of them before investigating the rest.

Seems like a very sensible approach and the trades it has identified would certainly be the ones our compliance people would want to look at.

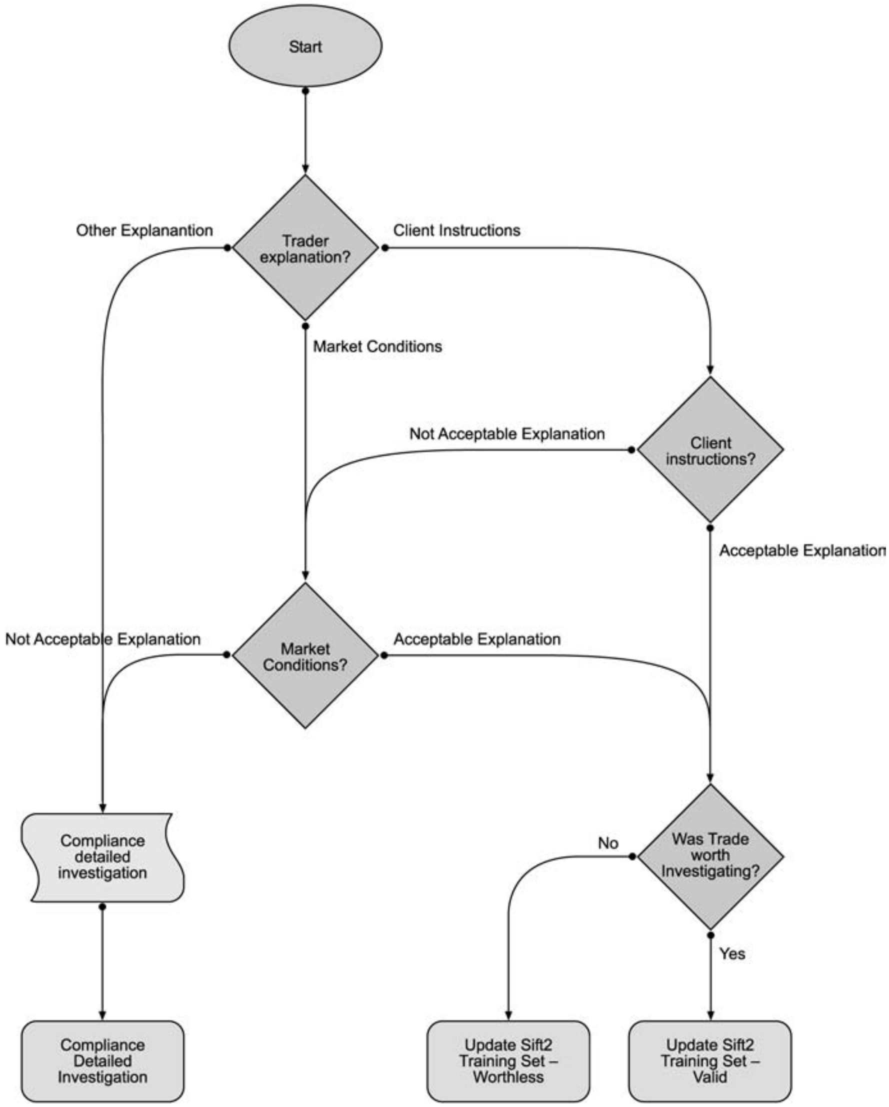

Figure 12.
Possible workflow management tool. Part 2 – investigation process

> The prototype workstation, with all the information about each trade and the drill down functionality is just the sort of tool we would use.

However, it is clear that the participant reactions so far are based on quick inspection, not statistical analysis. There does appear to be a need for an agreed set of "test trades" incorporating normal and anomalous trades, so that any system could use the test set for calibration and validation; i.e. does it spot trades that industry insiders would agree are anomalous?

How could an SVM/DAPR approach be used in practice?

In practice the SVM/DAPR approach would probably be combined with a workflow management system that could be integrated with the visualization and drill down tools shown in Figures 1 and 2. Initially the SVM/DAPR filter for outliers and the second filter for trades outside the bid/offer spread would be used to sift the anomalous trades from the regular ones. These trades would be combined with a small number of randomly selected trades as a control to produce a set of trades worthy of investigation (Figure 11).

The trades that are to be investigated could be referred to the trader for explanation – do client instructions, market conditions, the size of the trade or any other factors explain the execution price? (see Figure 12). This type of system would allow brokers to track the investigation process and provide an "audit trail" of compliance officers' work.

While this process might well be an adequate process, it might be improved by incorporating the results of investigations. For example, as these anomalous trades were investigated, over time a second sift could be constructed by building a second SVM/DAPR tool. This second sift would "learn" from the compliance team's investigations whether some suspicious trades were actually not worth investigation.

This research shows that dynamic anomaly and pattern response systems based on support vector machines appear capable of identifying trading anomalies at reasonable cost within usable timeframes. The choice for the wholesale financial industry is whether to move forward with SVM/DAPR as a proportionate automated response or to resist attempts to prove that trades are well executed.

References

Boser, B., Guyon, I. and Vapnik, V. (1992), "A training algorithm for optimal margin classifiers", *Fifth Annual Workshop on Computational Learning Theory (COLT-92), Pittsburgh, ACM*, pp. 144-52.

Burbidge, R. and Buxton, B. (2001), "An introduction to support vector machines for data mining", *Keynote YOR 12 (Young Operational Research 11th Conference)*, 28 March, University of Nottingham and Computer Science Department, University College London, available at: www.orsoc.org.uk/conf/previous/yor12/Keynote%20-%20Robert%20 Burbidge.doc

Cao, L.J. and Tay, F.E.H. (2001), "Financial forecasting using support vector machines", *Neural Computing Applications*, Vol. 10 No. 2, Springer-Verlag, New York, NYMay, pp. 184-92.

Cao, L.J. and Tay, F.E.H. (2003), "Support vector machine with adaptive parameters in financial time series forecasting", *IEEE Transactions on Neural Networks*, Vol. 14 No. 6, November, pp. 1506-18.

Huang, W., Nakamori, Y. and Wang, S.Y. (2005), "Forecasting stock market movement direction with support vector machine", *Computers and Operations Research*, Vol. 32 No. 10, pp. 2513-22, Elsevier, Barking.

Liddell, F.D.K. (1983), "Simplified exact analysis of case-referent studies; matched pairs; dichotomous exposure", *Journal of Epidemiology and Community Health*, Vol. 37 No. 1, pp. 82-4, March.

Mainelli, M. (2004), "The Michael Mainelli Column: finance looking fine, looking DAPR: the importance of dynamic anomaly and pattern response", *Balance Sheet*, Vol. 12 No. 5, pp. 56-9.

Mainelli, M. (2005), "The Michael Mainelli Column: competitive compliance: manage and automate, or die", *Journal of Risk Finance*, Vol. 6 No. 3, pp. 280-4.

Mainelli, M. and Yeandle, M. (2006), "Best execution compliance: new techniques for managing compliance risk", *Journal of Risk Management*, Vol. 7 No. 3, pp. 301-12.

Mainelli, M., Harris, I. and Helmore-Simpson, A. (2003), "The auditor's cross subsidy (statistical modeling of audit prices)", *Strategic Planning Society E-Newsletter, Article 1*, 19 June, also published as: anti-dumping measures and inflation accounting: calculating the non-audit subsidy, available at: www.mondaq.com.

Meyer, D., Leisch, F. and Hornik, K. (2002), "Benchmarking support vector machines", *Adaptive Information Systems and Modelling in Economics and Management Science Report Series, Number 78*, November, Vienna University of Economics and Business Administration.

Michie, D., Speigelhalter, D.J. and Taylor, C.C. (1994), *Machine Learning, Neural and Statistical Classification, Ellis Horwood*, out of print, available at: www.amsta.leeds.ac.uk/ ~ charles/statlog/

Vapnik, V.N. (1995), *The Nature of Statistical Learning Theory*, Springer-Verlag, New York, NY.

Vapnik, V.N. (1998), *Statistical Learning Theory*, John Wiley and Sons, New York, NY.

Further reading

Cristianini, N. and Shawe-Taylor, J. (2000), *An Introduction to Support Vector Machines and Other Kernel-based Learning Methods*, Cambridge University Press, Cambridge.