



Analyzing predictive performance of linear models on high-frequency currency exchange rates

Chanakya Serjam¹ · Akito Sakurai¹

Received: 31 July 2017 / Accepted: 10 May 2018 / Published online: 26 May 2018
© The Author(s) 2018

Abstract

We generate a large number of predictive models by applying linear kernel SVR to historical currency rates' bid data for three currency pairs obtained from high-frequency trading. The bid tick data are converted into equally spaced (1 min) data. Differences of price between the previous successive timeframes are used as features to predict the direction of movement of the price in the next timeframe. Different values for the number of training samples, number of features, and the length of the timeframes are used when learning the models. These models are used to conduct simulated currency trading in the year following the one in which the model was learned. Profits (sum of realized differences in best bid prices when order is executed), hit ratios, and number of trades executed using these models are recorded. The experiments indicate that while it is difficult to construct models using only historical data that consistently perform well, there are models that show good performance under certain pre-defined conditions, and that many of these models have an interesting property. Upon examining the parameters of these models, we discover that they have all negative coefficients and a negligibly small intercept, while having positive profits and good hit ratio. This suggests a simple yet effective trading strategy. Finally, we examine the historical data to find corroboration for the pattern suggested by the generated models and present the results.

Keywords Support vector regression (SVR) · Machine learning · Currency prediction · High-frequency limit order book

1 Introduction

Global financial markets have undergone a technological revolution in the past couple of decades. This has been made possible by the rapid advancements in various technical fields as well as major developments in the software and hardware in use. Many established exchanges have widely adopted

An earlier version of this research and paper was presented at the ACI-IDS 2017 conference at Kanazawa, Japan, in April 2017. The authors are grateful to the organizers of the conference and all the participants and reviewers who provided valuable comments and feedback.

This paper expands the training parameters used in the experiments to a much broader range, performs the experiments for another major currency pair (GB Pound/US Dollar), and investigates the historical data for the presence of the properties shown by the trained models.

✉ Chanakya Serjam
c.serjam.z3@keio.jp

Akito Sakurai
sakurai@ae.keio.ac.jp

¹ Graduate School of Science and Technology, Keio University, Yokohama 223-8522, Japan

electronic communication networks and trading systems [1–3]. As a result of the widespread acceptance and usage of the latest electronic systems in global financial markets, the processing time for tasks such as ordering or purchasing has gone down exponentially as compared to older traditional markets. Since lower processing time means lower overhead, the financial markets have a stake in pushing the processing time as low as possible. To achieve this, many financial marketplaces have been using high-frequency trading systems [4]. These systems keep human intervention (which is time-consuming and thus costly) to a bare minimum, and all the transactions are handled by computer algorithms to keep overhead such as time and cost as low as possible. High-frequency trading systems have been playing an increasingly vital role in trading (especially online trading). One major form of trading is currency trading or foreign exchange (forex for short). The forex market is certainly the largest, most liquid financial market in the world, dwarfing all other markets in size and volume of trading. However, it is also a very volatile market. As per a report from the Bank of International Settlements, the results from a recent survey [5] show that trading in foreign exchange markets averaged \$5.1 tril-

lion per day in just a single month (April) of 2016. Although this is down from an average of \$5.3 trillion per day in April of 2013, it is still a very voluminous market.

Traders investing in the currency markets are particularly interested in predicting the direction of movement of the price for the currency pair which they are looking to trade. If the price of the currency is about to go up, the trader will want to take the buy position, so he/she can sell the currency later at a higher price to turn a profit. If the price of the currency is about to go down, the trader will want to take the sell position. Later, the trader can buy the currency again for a lower price and turn a profit. Finally, the trader may assume a neutral position, i.e., neither buy nor sell. Therefore, the prediction task of a model trained for currency trading can have three outputs: buy, sell, or do nothing. The advent and widespread usage of high-frequency trading necessitates development and analysis of new trading strategies that can capture the short-term behavior of the markets. It is also very important to make an effort to understand the structure of the market under the influence of high-frequency trading.

In this paper, we conduct currency prediction experiments for Euro/US Dollar, British Pound/US Dollar, and US Dollar/Japanese Yen currency pairs using support vector machine for regression (SVR) [6,7], and examine the results to better understand the structure of currency trading in the forex market. Based on the forecast of the models, we perform simulated trading and record the profits or losses by comparing the predicted price movement with the actual price movement. We also examine the coefficient and intercept values and correlate them to the profit/loss and hit ratio metrics. The simulated trading is performed under some assumptions and defined pre-existing conditions that may not be representative of the real world but of an ideal scenario. Finally, we examine the historical data for the presence of properties exhibited by the models. Some interesting results are presented.

This paper is divided into the following major sections. Section 2 describes the background (previous research) and method of research. Section 3 describes the experimental setup and discusses the process in detail. Section 4 presents the results of the experiments and is used for analysis and discussion of the results. Finally, Sect. 5 presents a conclusion to the research and this paper.

2 Background and method of research

2.1 Previous research

While currency rates are volatile and prone to fluctuations, they have also been shown to be deterministically chaotic [8,9]. While this may be due to a number of factors, it is generally believed that historical data capture this behavior most concretely and effectively. Concurrently, historical data

usually become the primary input for any prediction model regardless of the technique used or the assumptions made.

A variety of techniques have been used for prediction tasks depending on the mathematical foundation or the value of specific model parameters. There has been considerable research [8–13] done on applying Artificial Neural Networks (ANNs) to forex forecasting. Deng et al. [14] and Deng and Sakurai [15] applied complex hybrid prediction techniques including Multiple Kernel Learning (MKL) and Genetic Algorithms (GA) to currency prediction and achieved good results. Kuo et al. [13] presented a decision support system for stock trading using GA-Based Fuzzy Neural Networks (GFNN) and ANNs. Another technique utilized for currency rates and financial timeseries prediction is Support Vector Machines [6,7], and it has also been applied successfully for high-frequency trading [16,17]. Studies [18,19] have shown that SVM-based models achieved on-par or better performance in forecasting of exchange rates or asset prices as compared to NN-based models for day trading.

While the techniques mentioned above show good results in prediction tasks, it is difficult to interpret the inner working of the models and how the prediction function generates the predictions. In addition, most of the techniques discussed above use dynamic training sets (using sliding window technique) to incorporate the latest data/information for making a prediction model. We were interested to know whether a model trained on a static training set can be used for prediction tasks far beyond the time horizon for which it is supposed to be valid. Therefore, due to combination of factors such as SVM techniques having good performance in financial forecasting tasks, the feasibility of linear models for understanding the prediction making process, and very little research available on using linear kernel SVR on static training set of historical data (only previous price differences) in high-frequency trading environment, we were motivated to perform this research.

2.2 Method of research

The primary aim of our research was to try and establish whether a linear model trained only on a static training set of historical data can have good predictive performance, and if so, to analyze the models to find out about the structure of the market. In our goal of analyzing the financial models which take historical data as input and produce relatively good performance, we planned to focus on the characteristics and structure of the model being generated. Hence, we decided on SVR with linear kernels to be the choice of technique for generating models, since it would be easier to analyze a linear model as the parameters would relate to real and observable data values. For further detailed reading and material on Support Vector Machines (SVM) and SVM for Regression (SVR), please refer to [6,7,16,17,20].

In high-frequency trading, the limit order book is updated every time there is a change in the bid or ask price or in case of other events such as a transaction being executed. These data are called tick data. The limit order book contains, among others, the timestamp (year/month/date and h/min/s), the best (highest) bid price, the bid volume, the best (lowest) ask price, and the ask volume. To study the timeseries properties of the price data, we only worked with the price data and eliminated the volume data. We also make use of only 1 price (bid price) rather than both the prices as there is not much qualitative difference in behaviors¹ between both prices. We also subjected the tick data to some pre-processing which included converting the tick data to equally spaced (1 min) data. Since the tick data are recorded every time there is a change in the order book, the data are unequally spaced and hence unsuitable for timeseries analysis. We wanted to check whether some patterns might emerge which can be learned by training models when the data are equally spaced. Converting the tick data to uniformly spaced data makes it easier to analyze as a timeseries.

In our experiments, we wanted to analyze whether there is a correlation between performance metrics such as profits or hit ratio and the initial parameters of the model such as size of the training set, the number of features to be used for prediction, and the length of timeframe (1, 2, 3 min, etc). Therefore, we trained models for many different values of these parameters. The models were trained on 1 year, and then used for validation on the data from the next year by performing simulated trading. This is to establish the predictive value of the models, since validating the models on the same year they were learned would not have yielded any information about the predictive performance of the models on new unseen data. Various performance metrics are observed and used for comparative analysis. Then, we examined the coefficients and intercepts of the models generated to look for some basic learning rule or pattern in the models. Finally, we analyze the historical data to see if the pattern suggested by the trained models is valid or not, and why a large number of models exhibit the same property.

3 Experimental setup

The currency rates data used in our experiments were acquired from ICAP. The experiments were performed on three different sets of currency pairs, the Euro/US Dollar data set, the GB Pound/US Dollar and the US Dollar/Japanese Yen data set. As previously mentioned, the original data sets contain the best bid and ask prices as well as the volumes.

¹ The initial experiments performed with both bid and ask price data for the sake of completeness revealed that the results using either price data are very similar.

The data sets are pre-processed to remove the volume data as well as the ask price data. Then, the tick data are converted to equally spaced (1 min) data which are the last tick data in the minute. Therefore, we have data sets that contain the date and the last price at each minute. The data sets used were from 2001 to 2015 and separated by year. Since the model is trained on the training set of the specified size extracted from 2 years (3 years in the case of GB Pound/US Dollar, since we need 3 years of minute data for GB Pound/US Dollar data to construct the required training set), and then used for prediction on the next year, the data of results for prediction analysis are from 2003 to 2015. For example, the models that were trained in the year 2001 and 2002 (2001 to 2003 for GB Pound/US Dollar) were used for prediction in the year 2003 (2004 for GB Pound/US Dollar); the models trained in 2002 and 2003 were used for prediction in 2004, and so on.

3.1 Parameters for training the models

- Number of features: The values used for the number of features were 1, 2, 3, 4, 5, and 6. Features used in our model are the difference of price between successive periods of time going back n periods from the current time (t). For instance, if the number of features is 1, it means that the model predicts the next output based on just one previous difference of price. Consequently, that model will have two parameters (since we are using linear kernel SVR), the coefficient and the intercept, and we extract those parameters to do a qualitative analysis of the model. If the number of features is n , the model predicts the next output based on n previous time frames and, therefore, the model will have $n + 1$ parameters.
- Length of timeframes: The lengths of timeframes (in minutes) used were 1, 2, 3, 4, 5, 7, 10, 20, 30, 40, 50, 60, and 70. These values were used to see if there is any correlation between the length of the timeframes and the performance metrics such as profits or hit ratio obtained. Although this could be extended to larger timeframes, we believe that it might not be fully reflective of the structure of high-frequency trading, where trading is very fast and timeframes are inherently small. We also considered that, in timeframes greater than 1 min, there may be multiple starting points from which the training set can begin. Therefore, we generate models for all the possible starting points (in minutes) within a timeframe and also average them.
- Size of training set: The values used for the number of training samples are 2000, 3000, 4000, 5000, 6000, 7000, 8000, 9000, and 10,000.

Models are generated for all possible combinations of these initial parameters.

Table 1 Summary statistics for the three currency pairs used in our experiments from the year 2001–2015

Currency pairs	Average no. of price quote updates per minute	Average no. of transactions (deals) per minute	Average bid-ask spreads for minute data
EURUSD	15.87	10.68	0.00021
GBPUSD	10.17	0.49	0.00087
USDJPY	13.90	6.73	0.02024

Table 2 Average ratios of positive changes vs negative changes for all 13 timeframes used in the experiments (table continued below)

Currency pairs	1 min	2 min	3 min	4 min	5 min	7 min	10 min
EURUSD	1.009	1.006	1.007	1.007	1.007	1.009	1.010
GBPUSD	1.020	1.013	1.011	1.009	1.008	1.007	1.006
USDJPY	1.016	1.014	1.013	1.013	1.014	1.016	1.017
Currency pairs	20 min	30 min	40 min	50 min	60 min	70 min	
EURUSD	1.011	1.013	1.011	1.012	1.011	1.011	
GBPUSD	1.008	1.010	1.011	1.012	1.014	1.014	
USDJPY	1.021	1.023	1.025	1.024	1.025	1.026	

3.2 Performance metrics

- **Hit ratio:** The hit ratio, also known as directional symmetry, is a measure of how many times the model predicted the change correctly. In other words, if the model predicts upward movement and the actual data used for validation confirm it, then it counts as a hit.
- **Profits:** Profits are obtained as a result of simulated trading based on the predictions of our models and are simply the sum of the realized differences in best bid prices when the orders are executed. If the price at the closing of a timeframe t is $price(t)$ and the prediction at the closing of the timeframe t is $pred(t)$, then profit is given as follows:

$$\text{Profit} = \sum [\text{price}(t + 1) - \text{price}(t)] \times \text{pred}(t). \quad (1)$$

For the Euro/US Dollar and the GB Pound/US Dollar currency pair, the profits were in US Dollars, and for the US Dollar/Japanese Yen currency pair, the profits were in Japanese Yen. It should be emphasized that the profits calculated in Eq. (1) are not representative of actual profits. In real-world trading, the concept of spread-crossing is an important and integral part of the profit calculation. Since we are working with only best bid prices, the spread does not factor into the equation. It is also important to point out that the bid-ask spread per trade is larger than the profits obtained per trade in most cases, and hence, profits calculated by Eq. (1) would not be positive if we did take the bid-ask spread into account.

For simulated trading, we put certain conditions in place. We assume that only 1 unit of the currency pair is being

traded. This was done under the assumption that a small transaction of 1 unit will not change or alter the market prices condition substantially and thus the following data set will not be disrupted. No fee is charged for transactions. In the real world, there is a small fee charged for every transaction, but we have chosen to ignore that to focus solely on the timeseries properties of currency trading.

In the simulated trading, a trade is counted when we have a change in the predicted direction of movement of the currency. Since we are only trading 1 unit, if, for instance, the prediction of direction is downward movement more than one times in a row, we do not execute or count those trades.

The summary statistics of the data for our experiments are displayed in Tables 1 and 2. Table 1 provides us metadata about average no. of price quote updates per minute, average no. of transactions (deals) per minute, and average bid-ask spreads observed for 1-min data. Table 2 shows the ratio of positive changes in best bid price vs. negative changes in the best bid price. Since all the ratios in Table 2 are slightly larger than 1, it implies that the number of positive changes in best bid prices has been slightly higher than the number of negative changes for all timeframes aggregated from 2001 to 2015. The next section discusses the results of the experiments.

4 Results and analysis

The results of the experiments consisted of the profits per year, the hit ratios, and the no. of trades executed over the period of a year using those models, as well as the intercept and coefficients of the models. Since the models were

grouped based on the number of features (1–6) used for the models, we calculated the average profits and hit ratios with respect to the length of timeframes and the size of training set (for each value of no. of features). This gave us four plots for each currency pair and gave insight into the performance of the models for different input parameters.

4.1 Performance metrics vs. length of timeframes

Figures 1, 2, and 3 below show the performance metrics (avg. hit ratio and avg. profits per year) as a function of the length of timeframe for all different values of number of features

(1–6) for the Euro/US Dollar pair, the GB Pound/US Dollar pair, and the US Dollar/Japanese Yen pair respectively.

It is interesting to note that, as the length of timeframe increases, the avg. hit ratio increases too irrespective of the no. of features, meaning an increase in the accuracy of trend prediction. However, at the same time, the profits from simulated trading go down as the length of timeframe increases. This is an interesting result, because normally profit would be expected to rise when hit ratio rises and vice versa. One reason for this might be that as the length of the timeframe increases, the no. of trades executed in our simulated trading decreases drastically. Thus, even if the hit ratio is higher,




Fig. 1 Plots for avg. hit ratio (left) and avg. profit per year (right) vs. length of timeframe for the Euro/US Dollar currency pair




Fig. 2 Plots for avg. hit ratio (left) and avg. profit per year (right) vs. length of timeframe for the GB Pound/US Dollar currency pair




Fig. 3 Plots for avg. hit ratio (left) and avg. profit per year (right) vs. length of timeframe for the US Dollar/Japanese Yen currency pair




Fig. 4 Plots for avg. hit ratio (left) and avg. profit per year (right) vs. training set size for the Euro/US Dollar currency pair

the number of trades executed might simply not be enough to generate profits comparable to shorter timeframes, which have lower hit ratio but a large number of executed trades, and thus more average profit per year.

In addition, we can see that fewer number of features results in higher hit ratio but lower profits on average.

4.2 Performance metrics vs. training set size

Figures 4, 5, and 6 show the performance metrics as a function of the size of the training set for all different values of number of features for the Euro/US Dollar pair, the GB Pound/US Dollar pair, and the US Dollar/JP Yen pair, respectively.

The plots show that there is an increase in both the hit ratio and the profits as the size of the training set increases. This might be because smaller training sets lead to over-fitting, whereas larger training sets can fine tune the parameters a bit better. On average, fewer number of features results in higher hit ratio and higher profits; although, in US Dollar/Japanese Yen (Fig. 6), lower profits for fewer number of features are observed.

4.3 Analyzing trained model parameters

While taking a cursory glance at our results, we noticed that a large number of models generated had similarities in the




Fig. 5 Plots for avg. hit ratio (left) and avg. profit per year (right) vs. training set size for the GB Pound/US Dollar currency pair




Fig. 6 Plots for avg. hit ratio (left) and avg. profit per year (right) vs. training set size for the US Dollar/Japanese Yen currency pair

correlation between the values of the intercept and the coefficients. These models had negligibly small intercept (which would not influence the predictions) as well as negative coefficients (although the number of models like this decreased as the no. of features, and thus the no. of coefficients, increased) while still giving good hit ratios and profits. We checked for the number of models that satisfied the condition of very small intercept, negative coefficients, and positive profit and hit ratio. The results for the Euro/US Dollar pair, the GB Pound/Japanese Yen pair, and the US Dollar/Japanese Yen pair are shown in Figs. 7, 8, and 9 respectively. The cases C1, C2, C3, and C4 are described as follows:

- Case 1 (C1): Absolute value of intercept < 0.1 , all coefficients < -10 (< 0 for US Dollar/Japanese Yen),² profits > 0 , and hit ratio $\geq 60\%$.
- Case 2 (C2): Absolute value of intercept < 0.1 , all coefficients < -10 (< 0 for US Dollar/Japanese Yen), profits > 0 , and hit ratio $\geq 50\%$ and $< 60\%$.

² The difference between the values of coefficients being checked for the US Dollar/Japanese Yen pair as compared to the other currency pairs is due to the difference in tick rate. US Dollar/Japanese Yen count the smallest tick at the second decimal place. The other two currencies count the smallest tick at the fourth decimal place.




Fig. 7 Stacked bar plots representing the percentage of models for cases C1–C4 for different values of no. of features on x -axis (Euro/US Dollar currency pair)




Fig. 9 Stacked bar plots representing the percentage of models for cases C1–C4 for different values of no. of features on x -axis (US Dollar/Japanese Yen currency pair)




Fig. 8 Stacked bar plots representing the percentage of models for cases C1–C4 for different values of no. of features on x -axis (GB Pound/US Dollar currency pair)

- Case 3 (C3): Absolute value of intercept < 0.1 , all coefficients < -10 (< 0 for US Dollar/Japanese Yen), profits > 0 , and hit ratio $< 50\%$.
- Case 4 (C4): Rest of the models (where not all coefficients are negative or absolute value of intercept > 0.1 , or profits < 0).

The stacked bar plots confirmed our initial observation that a large number of models had negative coefficients and negligible intercept values while giving profit and good hit ratio.

Thus, for models trained using linear SVR with a single feature, we can give a simple rule which states that the next prediction will be the opposite of the most recent (previous) movement direction. Concretely, if the previous trend is down, the model will predict up for the next change, and if the previous trend is up, the model will predict down for the next change. Using this simple trading rule, we get profit and good hit ratio in our simulated trading when using a single previous movement in direction of the price. This property is called *return reversal*. From the bar plots below, we can see that out of all the models with just one feature, a large percentage of models fall into case 1 of having positive profits and good hit ratio with negligible intercept and negative coefficients. This includes models from all the different timeframes used when training the models. The positive profits and high hit ratio suggest that the strategy may be viable under certain pre-defined circumstances irrespective of the timeframe used.

For models with two or more features, while case 1 is still a significant percentage of the total models, it decreases as the number of features increases. Since two or more previous difference in prices is being considered, it is possible that some of the features are negative, while others are positive. In this case, it is difficult to make a definitive statement about the presence of return reversal, as the condition of all negative coefficients is nullified. However, for n features, if all n features are the same sign, then we can see the next price

Table 3 Checking the return reversal property (in percentages rounded to two decimal places) for $t = 1, 5, 20,$ and 60 min for Euro/US Dollar historical bid price data

	1 min		5 min		20 min		60 min	
	- 1	+ 1	- 1	+ 1	- 1	+ 1	- 1	+ 1
- 1	46.97	53.03	47.38	52.62	47.00	53.00	47.03	52.97
+ 1	52.58	47.42	52.20	47.80	52.36	47.64	52.58	47.42
- 1, - 1	46.89	53.11	45.92	54.08	45.38	54.62	45.56	54.44
+ 1, + 1	53.34	46.66	54.03	45.97	53.78	46.22	53.86	46.14
- 1, - 1, - 1	46.18	53.82	45.00	55.00	44.07	55.93	43.95	56.05
+ 1, + 1, + 1	54.50	45.50	55.48	44.52	54.89	45.11	55.09	44.91

Table 4 Checking the return reversal property (in percentages rounded to two decimal places) for $t = 1, 5, 20,$ and 60 min for GB Pound/US Dollar historical bid price data

	1 min		5 min		20 min		60 min	
	- 1	+ 1	- 1	+ 1	- 1	+ 1	- 1	+ 1
- 1	45.90	54.10	46.69	53.31	46.67	53.33	46.51	53.49
+ 1	53.03	46.97	52.88	47.12	52.89	47.11	52.61	47.39
- 1, - 1	46.12	53.88	45.77	54.23	45.90	54.10	45.75	54.25
+ 1, + 1	53.34	46.66	54.18	45.82	54.01	45.99	53.74	46.26
- 1, - 1, - 1	45.67	54.33	44.96	55.04	45.16	54.84	45.20	54.80
+ 1, + 1, + 1	53.87	46.13	55.05	44.95	54.60	45.40	53.70	46.30

Table 5 Checking the return reversal property (in percentages rounded to two decimal places) for $t = 1, 5, 20,$ and 60 min for US Dollar/Japanese Yen historical bid price data

	1 min		5 min		20 min		60 min	
	- 1	+ 1	- 1	+ 1	- 1	+ 1	- 1	+ 1
- 1	46.13	53.87	47.10	52.90	46.91	53.09	46.76	53.24
+ 1	53.06	46.94	52.22	47.78	52.23	47.77	51.95	48.05
- 1, - 1	45.97	54.03	45.69	54.31	45.69	54.31	45.60	54.40
+ 1, + 1	53.86	46.14	53.85	46.15	53.73	46.27	52.71	47.29
- 1, - 1, - 1	45.12	54.88	44.37	55.63	44.60	55.40	44.39	55.61
+ 1, + 1, + 1	54.80	45.20	55.44	44.56	54.90	45.10	54.10	45.90

movement will be the opposite sign with a much higher probability³ irrespective of timeframe, as this would satisfy the models in case 1.

In the next sub-section, we take a look at the percentages of return reversal when all n features are the same sign for models with two or more features. We do this for different timeframes to see if the condition is still satisfied.

4.4 Checking historical data for occurrence of return reversal

We check for return reversal using 1, 2, and 3 features over sample timeframes of $t = 1, 5, 20,$ and 60 min for all three currency pairs. The reason which we chose to check for return reversal at those timeframes is because it provides a good spread from all the timeframes that we used to generate the

models. In case the change in price at the next step is 0, we look for the nearest non-zero value in the future. Only bid data are used, since we also used bid data in training the models.

Tables 3, 4, and 5 show the number of times (in percentages) the sign of the next value changes based on the previous consecutive opposite signs. The rows show the previous direction of movement of the price up to time t . - 1 s represent a negative change in price (the price goes down), whereas + 1 s represent a positive change in price (the price goes up). Concurrently, two or more consecutive - 1 s or + 1 s represent two or more such consecutive moves in the same direction. The columns show the probability of the following direction of movement of the price for time $(t + 1)$. The results are very consistent for all three currency pairs and for all the timeframes checked. One or more than one consecutive - 1 s is consistently followed by + 1 with a higher percentage or probability in all timeframes. Similarly, one or more than one consecutive + 1 s is consistently followed by

³ This does not mean that the actual movement will be of opposite sign. However, the accuracy of predicting the movement of the sign is greater than 50%, and cannot be described as purely chance.

$a - 1$ with a higher percentage in all timeframes. Thus, the probability of return reversal is always higher than that of the trend continuing irrespective of the timeframe or the length of the trend checked. This also helps to explain why a large number of models learned with linear kernel SVR, even for more number of features and for varied timeframes, showed properties of return reversal.

5 Conclusion and future work

In this paper, we conducted experiments to examine the performance of currency prediction models trained using linear kernel SVR on historical bid price data for high-frequency currency trading. We created models using various values for input parameters such as the length of training set, number of features, and length of timeframe for prediction. We also validated the results by performing simulated trading and recording the profits and hit ratio on next year's data and got good results. On examining the models, we found a simple rule that gave good results for models with single features, which is to predict opposite of the previous direction. This property is also known as *return reversal*. For models with two or more features, consecutive previous movements in the same direction will result in a higher probability of the next movement being in the opposite direction. Finally, we validated these findings by examining the historical data for occurrence of return reversal, and showed that the probability of the price movement changing directions is above the chance level, and that the property of return reversal holds true irrespective of the timeframe being used.

For future work, we plan to study models with more complex features, including technical indicators, and hope to find a trading strategy that incorporates return reversal but has even better performance. We also plan to do further analysis to establish the statistical significance of results obtained in this experiment. Finally, we also hope to create models which give a better prediction of return reversal based on several other features such as technical indicators generated from the price data.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Cont, R., Stoikov, S., Talreja, R.: A stochastic model for order book dynamics. *Oper. Res.* **58**(3), 549–563 (2010)
2. Parlour, C.A., Seppi, D.J.: Limit order market: a survey. *Handb. Financ. Intermed. Bank.* **5**, 63–95 (2008)
3. Bank of International Settlements: Triennial Central Bank survey of foreign exchange and derivatives market activity in 2007. <http://www.bis.org/publ/rpfx07t.htm>. Accessed 18 May 2018
4. Miller, R.S., Shorter, G.: High frequency trading: overview of recent developments (2016). <https://fas.org/sgp/crs/misc/R44443.pdf>. Accessed 18 May 2018
5. Bank of International Settlements: Triennial Central Bank survey of foreign exchange and OTC derivatives markets in 2016. <http://www.bis.org/publ/rpfx16.htm>. Accessed 18 May 2018
6. Cortes, C., Vapnik, V.: Support vector networks. *Mach. Learn.* **20**(3), 273–297 (1995)
7. Smola, A., Vapnik, V., et al.: Support vector regression machines. *Adv. Neural Inf. Process. Syst.* **9**, 155–161 (1996)
8. Hall, J.W.: Adaptive selection of U.S. stocks with neural nets. In: Deboeck, G.J. (ed.) *Trading on the Edge: Neural, Genetic and Fuzzy Systems for Chaotic Financial Markets*, pp. 45–65. Wiley, New York (1994)
9. Yao, J., Tan, C.L.: A case study on using neural networks to perform technical forecasting of forex. *Neurocomputing* **34**, 79–98 (2000)
10. Zimmerman, H., Neuneier, R., Grothmann, R.: Multi-agent modeling of multiple FX-markets by neural networks. *IEEE Trans. Neural Netw.* **12**(4), 735–743 (2001)
11. Zhang, G., Hu, M.Y.: Neural network forecasting of the British Pound/US Dollar exchange rate. *OMEGA Int. J. Manag. Sci.* **26**(4), 495–506 (1998)
12. Ni, H., Yin, H.: Exchange rate prediction using hybrid neural networks and technical indicators. *Neurocomputing* **72**(13–15), 2815–2823 (2009)
13. Kuo, R.J., Chen, C.H., et al.: An intelligent stock trading decision support system through integration of genetic algorithm based fuzzy neural network and artificial neural network. *Fuzzy Sets Syst.* **118**(1), 21–45 (2001)
14. Deng, S., Sakurai, A., Yoshiyama, K., Mitsubuchi, T.: Hybrid method of multiple kernel learning and genetic algorithm for forecasting short-term foreign exchange rates. *Comput. Econ.* **45**(1), 49–89 (2015)
15. Deng, S., Sakurai, A.: Integrated model of multiple kernel learning and differential evolution for EUR/USD trading. *Sci. World J.* **2014**(914641), 12 (2014)
16. Fletcher, T., Shawe Taylor, J.: Multiple kernel learning with Fisher kernels for high frequency currency prediction. *Comput. Econ.* **42**(2), 217–240 (2013)
17. Kercheval, A., Zhang, Y.: Modeling high-frequency limit order book dynamics with support vector machines. *Quant. Financ.* **15**(8), 1315–1329 (2015)
18. Tay, F.E.H., Cao, L.: Application of support vector machines in financial time series forecasting. *OMEGA Int. J. Manag. Sci.* **29**(4), 309–317 (2001)
19. Kim, K.: Financial time series forecasting using support vector machines. *Neurocomputing* **55**(1–2), 307–319 (2003)
20. Burges, C.J.C.: A tutorial on support vector machines for pattern recognition. *Data Min. Knowl. Discov.* **2**(2), 121–167 (1998)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.