

Deep Recurrent Modelling of Stationary Price Formation Using the Order Flow

Ye-Sheen Lim and Denise Gorse

University College London - Computer Science
Gower Street, London, WC1E 6BT

Abstract. In this paper we present a deep recurrent model based on the order flow for the stationary modelling of the directional movements of high-frequency prices. The formation of prices seen on the price chart of a stock or currency at any timescale are driven by the order flow, which is the microsecond streams of orders arriving at the exchange. In our experiments, we use high-frequency Bitcoin data to train our models. We train a deep recurrent neural network on the order flow to model the directional price movement given a sequence of order flow. We show that without any retraining, the model is temporally stable even as Bitcoin trading shifts into an extremely volatile "bubble trouble" period. The significance of the result is shown by benchmarking against existing state-of-the-art models for modelling price formation using deep learning. As such the secondary contribution of this paper is also a comparative study between the stationarity of existing deep models in literature.

1 Introduction

The main aim of this paper is the investigation of the stationary modelling of high-frequency price formation using deep learning approaches. Price formation is an important area in the study of market microstructure concerning the process by which asset prices form. When modelling price formation in practice, one is concerned with the *stationarity* of the model. Stationarity is the ability of a model to maintain prediction performance not just out-of-sample, but across a range of periods where the underlying process that generates the data undergoes drastic changes. The financial market is subjected to chaotic shift in regimes, and as a consequence, a model that is trained and tested in a particular period is not guaranteed to perform just as well in a future period if some unobservable underlying process of the financial market causes the statistical properties to drastically shift. Also, the stationary of price formation modelling is useful to academy as it is tied closely to the study of the financial market as complex dynamical system.

In this paper, we propose a deep recurrent model for modelling the *price formation* of Bitcoin using the *order flow*. The *order flow* is the microsecond timestamped sequence of events arriving to an exchange which causes the formation of prices we see on price charts. We formulate the *price formation* modelling problem as the modelling *directional price movements* of Bitcoin at *high-frequency*

as is common in quantitative finance literature [1, 2]. Bitcoin data is used in our experiment due to the ease of obtaining the high-frequency form of such data as opposed to the data for other financial assets. Also, we are able to obtain Bitcoin data covering the extremely volatile bubble period which is crucial for allowing us to study the stationarity of the models. We train our proposed order flow model on the Bitcoin data and show that our model is able to display the very desirable property of stationarity. We also implement and compare other deep learning models in literature in their stationarity.

2 Financial Background

Most, if not all, modern electronic stock or currency exchanges are driven by *limit order books* (LOBs). The electronic LOB is a platform that aggregates the quantities market participants desire to buy or sell at different prices. Most trading activities revolve around the lowest sell and highest buy prices. Readers are directed to [4] for a comprehensive introduction to LOBs.

Any exchange that uses a LOB is order-driven, such that any trader can submit *limit orders* (LO) to buy or sell a quantity of an asset at a specific limit price. If the order cannot be satisfied (at the specified limit price or better) on arrival to an exchange, then the LO is added to the LOB to be matched against subsequent orders arriving at the exchange. LOs in the LOB can also be cancelled at any time using a cancellation order (CO). Traders can also submit a *market order* (MO), which has no limit price and is always immediately executed at the best price in the LOB. The sequential stream of order book events is called the *order flow*.

Although there are many useful measures that can be computed from LOBs and order book events, of interest to this paper are the *mid-price*, *best bid*, *best ask* and the *relative price*. The *best bid* and *best ask* are the highest buying and lowest selling prices in the LOB respectively. The *mid-price* is the mean of the best ask and best bid, essentially the mid-point between the highest buying and lowest selling prices in the LOB. The *relative price* is the number of *ticks* between any two prices, where a *tick* is the lowest price increment or decrement allowed by the exchange.

3 Related Work

Theory-driven modelling of high-frequency (HF) price movements is an extensively researched topic. These approaches usually apply well defined stochastic models, chosen based on empirical analysis and market theories, for modelling HF price movements as a function of different measures of the LOB and order book events. Selected works in this area include [1, 11, 2]. The advantage of these approaches is their ability to simultaneously produce probabilistic forecasts of high-frequency price movements and confidently explain the predictions using the well-defined theory-driven models. However, major drawbacks include

reliance on parametric models, intractability of the models and the lack of generalisation power of the models.

Data-driven modelling of HF price movements is a relatively recent area of research. In machine learning literature, existing methods of modelling HF price movements differ mainly in how the covariates are treated, rather than the models themselves. Existing work with the most significant results relies on the power of recurrent neural networks to drive the price models. All these covariates are derived from the order flow, examples being the order imbalance and snapshots of the limit order book. Order imbalance is the difference between the available trading opportunities on the best bid and best ask, for a limit order book at a given point in time. The time series of order imbalance has been shown to be a good feature for forecasting high-frequency price movements [9]. Other authors have shown that the time series snapshot of the limit order book also holds strong predictors of HF price movements [10, 12]. The snapshot of the limit order book is the price and quantity of a given number of highest bid and lowest ask prices, at a given point in time. Augmenting the order book snapshot with the number of recent market orders (the arrival of which is guarantee to cause a price movement) has been shown to improve model performance [3]. However such order flow derived covariates are not necessarily ideal as will be shown later in the paper.

4 Methods

Our main goal is to model the probability distribution of the price movement y_i such that

$$P(y_i = j \mid \mathbf{h}_{i,T}^L, \mathbf{W}_j^D) = \frac{e^{z_j^D(\mathbf{h}_{i,T}^L, \mathbf{W}_j^D)}}{\sum_{k=0}^{K-1} e^{z_k^D(\mathbf{h}_{i,T}^L, \mathbf{W}_k^D)}} \quad , \quad (1)$$

where $j \in \{0, 1\}$ are $K = 2$ classes indicating the downward and upward price movements respectively, $\mathbf{h}_{i,T}^L$ is some learnt L -layer deep representation of order flow, and z_k^D is the output layer of a D -layer fully-connected neural network.

The representation $\mathbf{h}_{i,T}^L$ is the output of a deep L -layer recurrent neural network taken at the end of a length T order flow:

$$\mathbf{h}_{i,T}^l = \begin{cases} h(\mathbf{h}_{i,T}^{l-1}, \mathbf{h}_{i,T-1}^l, \Theta^l) & \text{if } 1 < l \leq L \\ h(\mathbf{x}_{i,T}, \mathbf{h}_{i,T-1}^l, \Theta^l) & \text{if } l = 1 \end{cases} \quad , \quad (2)$$

where $h(\cdot)$ is a function implementing a recurrent neural network with long short-term memory (LSTM) cells, Θ^l are LSTM parameters to be fitted, and $\mathbf{x}_{i,T}$ will be soon addressed. Since LSTM cells are commonly implemented RNN components in the literature, we will not discuss their architecture here and direct readers instead to [5]. The variable \mathbf{x}_i is our covariate and can be defined as a sequence of T orders $\mathbf{x}_{i,t}$ as follows:

$$\mathbf{x}_i = (\mathbf{x}_{i,1}, \mathbf{x}_{i,2}, \mathbf{x}_{i,3}, \dots, \mathbf{x}_{i,T}) \quad , \quad (3)$$

$$\mathbf{x}_{i,t} = (\delta_{i,t}, h_{i,t}, \omega_{i,t}, q_{i,t}, \alpha_{i,t}, \pi_{i,t}) \quad , \quad (4)$$

where

- $\delta_{i,t} \in \mathbb{N}$ is the number of milliseconds between the arrival of $x_{i,t-1}$ and $x_{i,t}$
- $h_{i,t} \in \mathbb{N}$ is the hour of the arrival of $x_{i,t}$ according its timestamp
- $\omega_{i,t} \in \mathbb{R}$ is the size of the order $x_{i,t}$, such that $\omega_{i,t} > 0$
- $q_{i,t} \in \{1, 2, 3\}$ is the categorical variable for $x_{i,t}$ being a limit order, market order or cancellation order
- $\alpha_{i,t} \in \{1, 2\}$ is the categorical variable for $x_{i,t}$ being a buy or sell order
- $\pi_{i,t} \in \mathbb{N}$, such that $\pi_{i,t} \geq 0$, is the relative price of the order $x_{i,t}$ (if $\alpha_{i,t} = 1$ then the price is relative to the highest buy price in the LOB, and if $\alpha_{i,t} = 2$ then it is relative to the lowest sell price)

For each order $x_{i,t}$ all non-ordinal categorical covariates are embedded into a multidimensional continuous space before feeding into the inputs of the RNNs. Then, the parameters of the model \mathbf{W}_k^d and Θ^l , as well as the weights and bias terms in the embedding layers for all covariates, can be fitted using any variant of the stochastic gradient descent optimisation algorithms by minimising the negative log-likelihood:

$$\mathcal{L}(\mathbf{y}) = -\frac{1}{N} \sum_{i=1}^N \sum_{j=0}^{K-1} \mathbb{I}_{y_i=j} \log p(j) \quad (5)$$

where \mathbb{I} is the identity function that is equal to one if $y_i = j$ and is zero otherwise, K is the number of classes, and N is the size of our dataset.

The performance of the model will be evaluated using the Matthews correlation coefficient (MCC) [8]. We choose this metric as it has a very intuitive interpretation, it handles imbalanced classes naturally. For binary classification, the metric lies in the range $(-1, 1)$ where 1 indicates a perfect classifier, -1 indicates a completely wrong classifier and 0 means the classifier is doing no better than making random predictions (making this measure very useful in the context of quantitative trading). As it summarises the confusion matrix into one balanced and intuitively interpretable measure, it allows us to perform concise and extensive comparisons without needing to delve into the relative contributions of different elements of the confusion matrix.

We benchmark the performance of our order flow model against state-of-the-art models in literature literature for predicting high-frequency price movements, which were discussed in previous sections. These models are similar deep RNN models but with the input covariates defined as follows:

- Order Book Snapshot:

$$\text{snap}_i = (\text{snap}_{i,1}, \text{snap}_{i,2}, \dots, \text{snap}_{i,T}) \quad , \quad (6)$$

$$\text{snap}_{i,t} = (b_{i,t}^1, b_{i,t}^2, \dots, b_{i,t}^S, s_{i,t}^1, s_{i,t}^2, \dots, s_{i,t}^S) \quad , \quad (7)$$

where $b_{i,t}^j = (\pi_{i,t}^j, \omega_{i,t}^j)$ is the tuple of price π and volume ω at the j 'th highest buy price, $s_{i,t}^j = (\phi_{i,t}^j, \kappa_{i,t}^j)$ is the tuple of price ϕ and volume κ at the j 'th lowest sell price, and S is the number of highest buy or lowest sell prices we are considering in the snapshot

- Order Book Snapshot with Market Order (MO) Rates:

$$agg_i = (agg_{i,1}, agg_{i,2}, \dots, agg_{i,T}) \quad , \quad (8)$$

$$agg_{i,t} = (b_{i,t}^1, b_{i,t}^2, \dots, b_{i,t}^S, s_{i,t}^1, s_{i,t}^2, \dots, s_{i,t}^S, \alpha_{i,t}^b, \alpha_{i,t}^s) \quad , \quad (9)$$

where $b_{i,t}^j$ and $s_{i,t}^j$ are the same as in the order book snapshot without MO rates, $\alpha_{i,t}^b$ is the number of buy MOs arriving among the past T order book events divided by the number of orders that makes up the volume in the highest bid price, and $\alpha_{i,t}^s$ is the same but for sell MO orders to the lowest sell price

- Bid-Ask Order Imbalance:

$$imba_i = (imba_{i,1}, imba_{i,2}, \dots, imba_{i,T}) \quad , \quad (10)$$

$$imba_{i,t} = \kappa_{i,t}^1 - \phi_{i,t}^1 \quad , \quad (11)$$

where $\kappa_{i,t}^1$ and $\phi_{i,t}^1$ are as defined above

5 Experimental Set-Up

The dataset for the experiments was obtained from Coinbase, a digital currency exchange. Through a websocket feed provided by the exchange, we log the real-time message stream containing trades and orders market data updates in the form of order flow for BTC-USD and BTC-EUR from 4 Nov 2017 to 29 Jan 2018. Since these are raw JSON messages, the dataset for training the model cannot be obtained directly from the messages. To build the datasets, we had to reconstruct the limit order book by sequentially processing the JSON messages.

During the order book reconstruction, we build the dataset in real-time using the following method. Before we begin building the dataset, we have to "warm up" the order book using messages from 4 Nov 2017 - 5 Nov 2017. This needs to be done because we are starting from an empty book that needs to be sufficiently populated for us to extract sensible data. Now we begin tracking the mid-price. If the mid-price moves after an order $x_{i,T+1}$ arrives, then the mid-price is stored as the target variable class $y_i \in \{0, 1\}$, for downward or upward mid-price movements respectively. We ignore any events that do not move the price on arrival since we are only interested in predicting up or down price movements. After a mid-price change is registered, the covariates \mathbf{x}_i associated with a directional price change are taken to be the past T orders *excluding* order $x_{i,T+1}$. At the same time, the covariates for the benchmark models as described in the previous section are also computed, with great care taken to ensure the covariates of all models are sharing the same target variable for any given datapoint.

Each dataset is then split into training, validation and test sets by dates to avoid any look-ahead bias. About 1.1 million datapoints between 6 Nov 2017 and 16 Nov 2017 are taken as the training set. Cross-validation and early-stopping checks are performed on a set taken from between 17 Nov 2017 to 20 Nov 2017 containing about 0.55 million datapoints. The rest is kept strictly for testing and is never seen by the model aside from the final testing phase, giving us about 7.3 million test points in total.

We set up the order flow and benchmark models as described in Methods. Then we fit the parameters of individual models using the Adam optimisation algorithm with dropout on all non-recurrent connections. Cross-validation is performed with Bayesian hyperparameter tuning to find the number of recurrent layers, the LSTM state size, the number of dense layer in the output, the size of each output dense layers, and the embedding dimensions. Embedding is not used for the benchmark models since they do not contain categorical variables. We then make predictions on the test sets. The results are then grouped by date and we compute the Matthews correlation coefficient (MCC) for each date.

6 Results

In this section, for readability we use intuitive abbreviations for the names of the features: i) order *flow*, ii) order book *snapshot*, iii) order book snapshot with aggressors (market order rates), iv) bid-ask volume *imbalance*.

6.1 Comparison of Model Performances

Let us first evaluate the performance of the various models trained. Figure 1 shows the MCC evaluated at each day for models trained and tested on the BTC-USD dataset. We observe that throughout the test period, the performance of models trained on the order flow outperforms those of the benchmark models. Although not presented, we also performed paired Student t-tests on the null hypothesis that there is no difference between the test results of *flow* and that of benchmark models individually. In each test, the null hypothesis is rejected with very high confidence interval.

Similarly, Figure 2 shows the test performance for models trained and tested on BTC-EUR. We are able to visually verify that models trained on order flow outperform the benchmark models. Once again paired Student t-tests reject with high confidence the null hypotheses that *flow* test results are no different from the benchmark models.

6.2 Analysis of Model Stationarity

The range of dates in our test period covers the climax of the Bitcoin bubble where the price of Bitcoin rapidly peaks to an all-time-high and subsequently bursts [6]. We can see in Figure 3 that compared to the training and validation




Fig. 1. MCC plots for models trained and tested on BTC-USD




Fig. 2. MCC plots for models trained and tested on BTC-EUR

period, our test period for BTC-USD (and in fact for BTC-EUR also) corresponds to a shift in regime characterised by more volatile price changes and much higher trading activity. With this test period, we can evaluate the stationarity of our proposed order flow model and the benchmark model.

From Figures 1 and 2, we can in both cases (BTC-USD, BTC-EUR) visually observe the stationarity of the prediction performance of the model trained on order flow throughout the test set. This implies that the representation learnt from the order flow is transferable from a non-volatile to an extremely volatile period, suggesting that the learnt representation encodes some sort of temporally universal information about Bitcoin price movements. On the other hand, we can visually observe that the benchmark models that were trained on features computed from the order flow struggle to maintain performance in the volatile period. To scientifically evaluate and statistically compare the stationarity of the model trained on order flow and benchmark models, we fit a linear regression model on the test results of each model. This will then allow us to compare the rate at which the performance of each model degrades over time.

Specifically, we fit a simple linear regression on the MCC over the test period dates. Table 1 shows the corresponding slope coefficients and p-values for each




Fig. 3. Plot of: i) volume of trading activity (Trade Volume), ii) BTC-USD 1-day lagged difference prices (Price Change), the shaded area being our test period 21 Nov 2017 - 29 Jan 2018

model trained and tested on the BTC-USD and BTC-EUR datasets respectively. Although the model trained on order flow has negative slope coefficients for all datasets, implying some degradation in performance over time, the p-values reveal that the coefficient is not statistically significant. However, for the benchmark models *agg*, *snap*, *imba*, we can very confidently reject the null hypotheses that the slopes are zero, meaning that we can use the negative coefficients as strong evidence for performance degradation over time.

Dataset	Model	Coeff	p-value
BTC-USD	flow	$-2.41e^{-4}$	$5.87e^{-2}$
	agg	$-2.76e^{-3}$	$1.27e^{-13}$
	snap	$-1.63e^{-3}$	$8.27e^{-4}$
	imba	$-1.46e^{-3}$	$1.07e^{-5}$
BTC-EUR	flow	$-5.67e^{-5}$	$7.50e^{-1}$
	agg	$-2.02e^{-3}$	$1.27e^{-8}$
	snap	$-2.01e^{-3}$	$1.19e^{-5}$
	imba	$-1.01e^{-3}$	$1.56e^{-3}$

Table 1. The table shows the slope coefficients and p-values of the MCC regressed on dates in the test period for individual models that are: i) trained and tested on BTC-USD, ii) trained and tested on BTC-EUR

7 Analysis of Model Universality

Although we set out to study stationarity in the model, it is in addition possible to show that the representations learnt from the order flow exhibit a hint of the very valuable property of universality [10], the ability to learn market structures which to some degree generalise across asset classes. Table 2 shows the drop in

performance on the out of sample test set, when training on one currency pair and testing on another, is considerably less when using the order flow model than the benchmark models, demonstrating the above-mentioned hint of universality. This innate ability to generalise is most evident when training on BTC-USD compared to BTC-EUR. This is likely due to the different volumes traded: the trading volumes of BTC-USD and BTC-EUR between the start of the training period (6 Nov 2017) and the end of the test period (29 Jan 2018) are, respectively, $151.6e^9$ and $20.7e^9$ – when an asset is heavily traded there are more activities at the order level, resulting in a richer order flow and hence a richer dataset that helps to avoid overfitting.

Trained on:	BTC-USD	BTC-EUR
Tested on:	BTC-EUR	BTC-USD
flow	9.299	19.606
agg	59.857	74.368
snap	91.017	99.619
imba	87.504	83.776

Table 2. For each model, the presentation tabulates the mean percentage drop (%) in test MCC between training and testing on a single currency pair, and training on a given currency pair and testing on the other currency pairs

8 Discussion

We have presented an order flow model for the deep stationary modelling the directional price movements of Bitcoin using the order flow (which is the raw market data). We showed that the model is able to partially achieve a temporally universal representation of the price formation process such that even when the statistical behaviour of the market changes due to a high-impact exogenous "Black Swan" event, the model remains relatively unaffected. We also show that the stationarity performance of our proposed order flow model are substantially better than the benchmark deep models obtained from existing literature. A secondary analysis of the results also hints at the universality property of our proposed model, and this is also benchmarked against deep models in existing literature.

For future work, since our predictions give encouragingly high MCC values, it will be of interest to apply black-box feature explainers such as Shapley Additive Explanations [7] to address the interpretation issue of these data-driven models and understand exactly what it is that drives price formation across BTC-USD and BTC-EUR, and (in future work) other cryptocurrency pairs. This would provide a data-driven view of the market microstructural behavior of cryptocurrencies. Also of interest is what we can learn about the cryptocurrency market microstructure from analysis of the embeddings of categorical features of the order flow.

Finally, we might ask to what degree the current models could be useful as pre-trained networks for use in related problems. In our analysis of universality, we show that it is possible to use the order flow model for the prediction of directional price movements for cryptocurrency pairs other than those on which the models were trained. However, we might go further and ask if the learnt representations, with their displayed degree of stationarity, could be used to initialise training, and by this route lead to faster convergence, for directional forecasting problems outside of the cryptocurrency context. Perhaps most interestingly, we could investigate whether models trained on high frequency data have learnt anything sufficiently general about the market microstructure so as to be relevant to a different (lower) timescale, where less data would be available and where a pre-trained learner would potentially be very valuable.

References

1. Bacry, E., Muzy, J.F.: Hawkes model for price and trades high-frequency dynamics. *Quantitative Finance* **14**(7), 1147–1166 (2014)
2. Cont, R., Stoikov, S., Talreja, R.: A stochastic model for order book dynamics. *Operations Research* **58**(3), 549–563 (2010)
3. Dixon, M.: Sequence classification of the limit order book using recurrent neural networks. *Journal of Computational Science* **24**, 277–286 (2018)
4. Gould, M.D., Porter, M.A., Williams, S., McDonald, M., Fenn, D.J., Howison, S.D.: Limit order books. *Quantitative Finance* **13**(11), 1709–1742 (2013)
5. Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Computation* **9**(8), 1735–1780 (1997)
6. Kreuser, J.L., Sornette, D.: Bitcoin bubble trouble. Forthcoming in *Wilmott Magazine* pp. 18–24 (2018)
7. Lundberg, S.M., Lee, S.I.: A unified approach to interpreting model predictions. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) *Advances in Neural Information Processing Systems* 30, pp. 4765–4774. Curran Associates, Inc. (2017), <http://papers.nips.cc/paper/7062-a-unified-approach-to-interpreting-model-predictions.pdf>
8. Powers, D.M.W.: Evaluation: From precision, recall and f-measure to roc., informedness, markedness & correlation. *Journal of Machine Learning Technologies* **2**(1), 37–63 (2011)
9. Shen, D.: Order imbalance based strategy in high frequency trading. Ph.D. thesis, Oxford University (2015)
10. Sirignano, J., Cont, R.: Universal features of price formation in financial markets: perspectives from deep learning. *Quantitative Finance* pp. 1–11 (2019)
11. Toke, I.M.: An introduction to hawkes processes with applications to finance. *Lectures Notes from Ecole Centrale Paris, BNP Paribas Chair of Quantitative Finance* **193** (2011)
12. Tsantekidis, A., Passalis, N., Tefas, A., Kannianen, J., Gabbouj, M., Iosifidis, A.: Using deep learning to detect price change indications in financial markets. In: *2017 25th European Signal Processing Conference (EUSIPCO)*. pp. 2511–2515. IEEE (2017)