REGIME CHANGE AND TREND PREDICTION FOR BITCOIN TIME SERIES DATA

Osamu Kodama,[1] Lukáš Pichl,[2] Taisei Kaizoji[3]

Abstract: Bitcoin time series dataset recording individual transactions denominated in Euro at the COINBASE market between April 23, 2015 and August 15, 2016 is analyzed. Markov switching model is applied to classify the regions of varying volatility represented by three hidden state regimes using univariate autoregressive model and dependent mixture model. Causality extraction and price prediction of daily BTCEUR exchange rates is performed by means of a recurrent neural network using the standard Elman model. Strong correlations is found between the normalized mean squared error of the Elman network (out-of-sample 5-day-ahead prediction) and the realized volatility (sum of minute returns squared throughout the trading day). The present approach is calibrated using simulated regime change in standard econometric models. Our results clearly demonstrate the applicability of recurrent neural networks to causality extraction even in the case of highly volatile cryptocurrency exchange rate time series data.

UDC Classification: 004.8, 33; **DOI:** http://dx.doi.org/10.12955/cbup.v5.954

Keywords: Bitcoin, BTC, Elman model, Hidden Markov Model, HMM, recurrent neural network.

Introduction

Bitcoin is a cryptocurrency released as an open-source software in 2009, which represents a transaction payment system as well as a sort of digital commodity (Bohme et al., 2016). The bitcoin market capitalization as of early 2017 has reached USD 20 billion (CoinDesk, 2017). Free of any interventions from regulatory authorities, such as central banks, the distributed block chain system on which Bitcoin is based meets varying levels of demand for transaction settlement, and the Bitcoin exchange rate series to major currencies such as USD, EUR or GBP are known to highly fluctuate – it is not uncommon that gains or losses in tens of percent occur within a week, if not during a single day.

There exist various Bitcoin exchange markets, such as BitBay, Btcde, Kraken, LocalBtc, or Rock for EUR currency, to name just a few of the currently active BTCEUR exchanges. The highly volatile nature of the exchange rate represents an ideal environment for the study of the extreme events in the field of financial time series. Prediction of extreme events is a key issue not only in economics, but also in climatology, geosciences, civil engineering, space technology, etc. In spite of its importance, the topic is rather under-studied, in our opinion.

In this work, we explore the applicability of computational intelligence methods from financial analysis to the series of Bitcoin exchange rates (data shown in Fig. 1). Since the Bitcoin price process is not stationary but exhibits an appreciation trend, we transform the time series data to the logarithmic returns. If the absolute value of the log return is large, it corresponds to an extreme event (bullish or bearish, based on the sign). Next we adopt the Hidden Markov Model to categorize the market regime into 3 modes: stable, intermediate, and volatile. This approach is excellent in ex-post analysis of the data, however lacks in the predictive power for future trend prediction. Consequently, we add the realized volatility as an intraday indicator of market stability, and develop a recurrent neural network configuration, which uses the past log return history in a moving window to predict the next week's log return behavior. Since the market process mixes both deterministic and stochastic modes, it is not a priori clear where the limit of the predictive power of the recurrent neural network is bound. This paper shows as the principal result that the mean squared error of the prediction is only limited by the level of the realized volatility.

This paper is organized as follows. Following the literature review in the next section, in Section 3 we explain the dataset and outline the methods of its analysis. Section 4 wraps up our results and discussions, which are followed by the concluding section.

[1] International Christian University, Mitaka, Tokyo, Japan, deep86.osa@gmail.com
[2] International Christian University, Mitaka, Tokyo, Japan, lukas@icu.ac.jp
[3] International Christian University, Mitaka, Tokyo, Japan, kaizoji@icu.ac.jp

Figure 1: Time series of BTC exchange rate to EUR at COIN-Base market (upper panel). Daily log return and realized volatility (on minute scale) for the same trading period (lower panel).

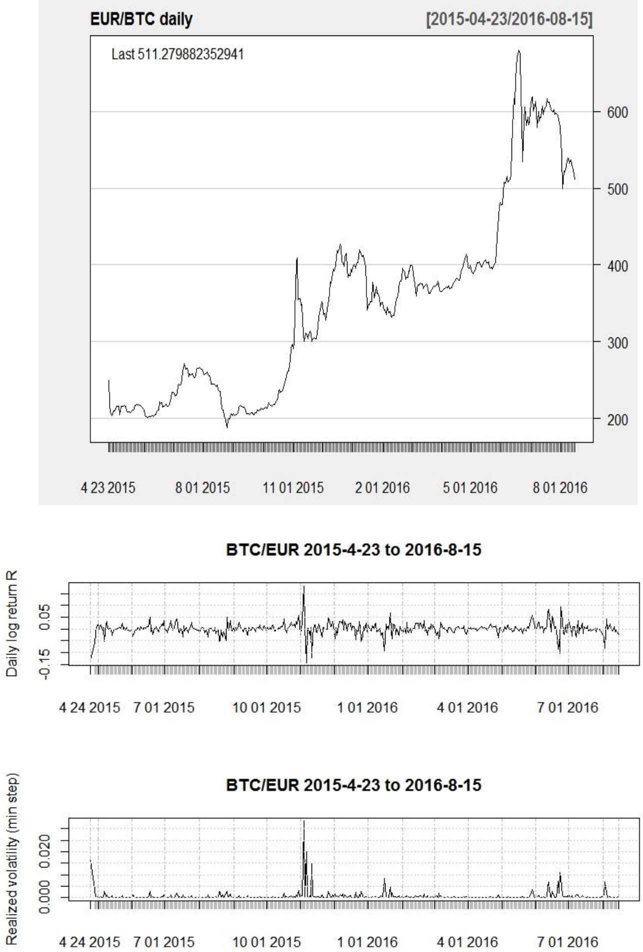

Source: Authors (data retrieved from COINBASE market)

Literature Review

Among the notable attempts to model the prediction of extreme events in a systematic way are those of Hallerberg et al., (2008) assessing under what circumstances the extreme events may be more predictable the bigger they are, or the recent work by Franzke (2012) who develops a nonlinear stochastic-dynamical model. In the economic context, extreme events mean a bubble formation or a bubble burst, and their precursors are of vital importance in risk management. To extract the causal extent (deterministic segment) buried in the noisy data, various techniques have been proposed, for instance recurrent neural network with memory feedback (Elman, 1990) or support vector machines (Cortes and Vapnik, 1995). A survey of recent methods can be found in the work of Akansu et al. (2016). Binary classifiers separating the upward and downward trend (positive or negative sign of logarithmic return), which easily evaluate against the dataset in terms of hit ratios (precision of binary classifier output), are common.

Data and Methods

The dataset of BTCEUR tick data between 2015-4-23 and 2016-8-5 contains 809,489 records in 478 trading days transacted at COINBASE market. All data analyses were performed on a Dell PowerEdge T420 server, with 2 Intel Xeon E5-2407 2.2 GHz processors and 8 cores, running the GNU/Linux Fedora 23 operating system.

Figure 1 shows the EUR exchange rate data of BTC, including the logarithmic returns on daily scale ($R_t=\log(P_t/P_{t-1})$) and the realized volatility (RV), which is computed as the sum of the logarithmic returns squared on minute sampling scale during the trading day. The higher the realized volatility, the bigger are the uncertainty and spread of trading values typically observed during the day. The probability density of these two quantities is depicted in Fig. 2.

Figure 2:Density functions for the log return and realized volatility from Fig. 1

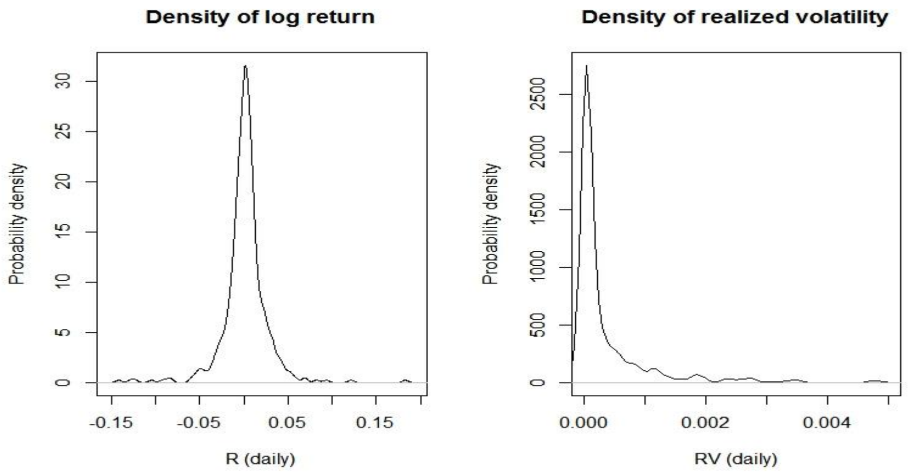

Source: Authors

The first approach to classification of the market trend regime is the state switching model based on Hidden Markov Model (HMM). The observable quanity is the sequence of the logarithmic returns. We set the number of hidden states to equal to 3, that is a stable, non-volatilite regime, highly volatile regime, and an intermediate state in between the two. The extraction of the underlying state probabilities is performed using the R-packages depmixS4 and MSwM (https://cran.r-project.org/). The theoretical approach closely follows that of Kirikos (2000). While the categorization results are quite reasonable, which can be seen in Fig. 3, namely the 3-states are correctly classified in regard to the magnitude of the logarithmic return, the predictive power of this approach is quite low (i.e. if we assume that the state for the next day is the one for the current day computed by the HMM model). We have to therefore resort to a more powerful causality extraction model. See (Gyorfi et al., 2012) for a list of possible candidates in the field of machine learning algorithms.

Figure 3: 3-state HMM for data of Fig. (1) (State #1: stable, #2 intermediate, #3 volatile)

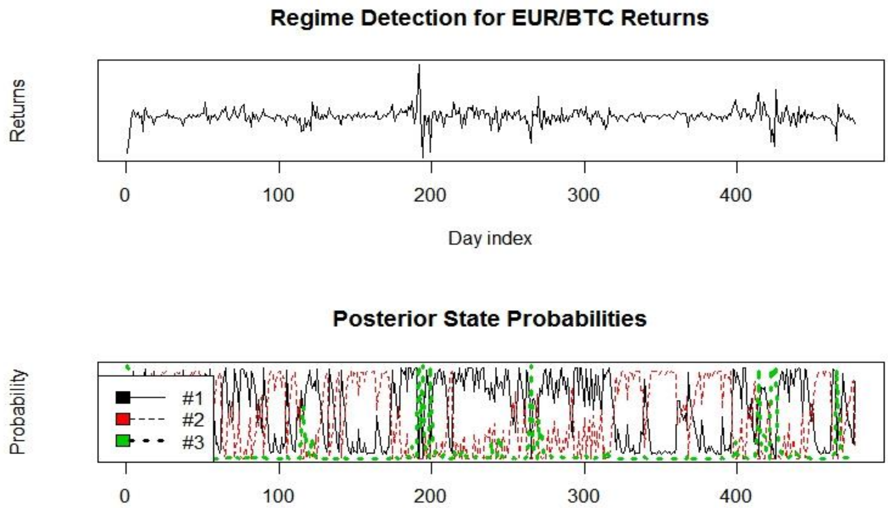

Source: Authors

In order to account for the time series correlations accurately, we have decided to apply the recurrent neural network (RNN) in Elman configuration (Elman, 1990). We have adopted two different topologies, one with a single hidden layer (8 units), and another one with two hidden layers (each from 5 to 8 units). The contents of the hidden units are fed as an extra input to the network in a feed-back loop, thus implementing the concept of the state memory. The implementation is that of the Stuttgart Neural Network Simulator, available as an R-package RSNNS. The main research question is whether the RNN model can extract the causal extent of the time series, and how it relates to the series of realized volatility, which are a measure of the stochastic content of the time series data.

Results and Discussions

The main result of the study is given in Fig. 4 for the configuration with two hidden layers, each containing 8 units. Thus there are two feed-back loops in the network configuration of RSNNS. The following procedure has been used: the RNN model is trained on the past 40-days of input data, and evaluated using a 5-day ahead prediction. The possible error outcome is therefore the precision set from {0.0, 0.2, 0.4, 0.6, 0.8 and 1.0}. The subset of {0.0, 0.2, 0.4} – that is majority-winning correct classification of the binary trend (sign of the logarithmic return) – has the frequency of 70.5 %, attesting to the causal content of the series. More importantly, we evaluate the RNN model using the standard measure of the Mean Squared Error of the logarithmic return for the 5 predicted values. Because of the 40-day moving window, this indicator is unavailable for the first 40 days (zero-level flat start in Fig. 4), and the curve is further 5 days shifted, in order to compare with the realized data of the 5-day ahead prediction window. In addition, normalization is used, i.e. the neural network receives data standardized to zero-mean and unit standard deviation; these units are applied to the evaluation of the MSE. The resulting curve (thick line in Fig. 4, red color in online version) shows an almost perfect coincidence, enveloping properly scaled graph of the realized volatility. Since the realized volatility measures the noise in the system, i.e. the unpredictable component of the time series, we can infer, based on Fig. 4, that the underlying deterministic content, i.e. the causal behavior mode, has been extracted properly using the RNN model. The results in Fig. 4 practically do not depend on RNN topology within the limit described above.

To further substantiate the scope of the validity of the RNN model, we have simulated econometric series of ARMA model (auto-regressive moving average model) with a single propensity parameter on the scale of 0.1 to 0.9. The RNN model reacts to the model discontinuity by the increase in the MSE, thus detecting the regime change. The bigger the parameter change, the better the chance is that the model change is discovered by means of an MSE increase.

Figure 4: Learning power of the RNN model: comparison of MSE vs. realized volatility

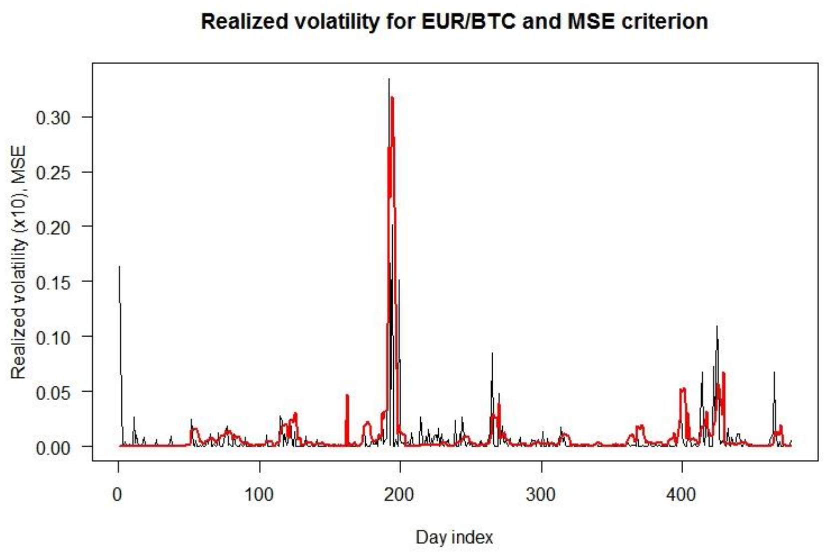

Source: Authors

Conclusion

This work has established the applicability of the Recurrent Neural Network (RNN model) to the time series of Bitcoin exchange rates denominated in EUR currency. We have derived the average time series for each day in a 478-day-long sample of tick data from the COINBASE market, computed the realized volatility, and analyzed the causal extent of the daily time series using both the standard HMM model and the RNN model. The HMM is a poor predictor of the regime changes as well as market trends. The RNN model, on the other hand, showed a predictive power related to the spikes of the MSE value. Using the realized volatility, we could see that the model performs up to the theoretical bounds of its applicability, capturing the full scope of the deterministic contents, with MSE therefore closely following the distribution of the stochastic error given by the RV distribution. The presented results show a good agreement of the MSE curve and RV distribution. This agreement is the better the higher the volatility spike is, which would conform to the model "the bigger the extreme event is, the better predicted it can be." Nevertheless, the presented results are still confined to a relatively short period and a single market, and thus further investigation is required to make a more general conclusion. We also plan to study the effect of market-making information using open access texts (cf. Kim et al., 2016). The present findings may also be useful in considerations of the design of future cryptocurrencies other than Bitcoin (Extance, 2015). The presented work is also relevant and interesting due to the limited number of available data analysis papers on the Bitcoin subject, although the situation has improved recently (Houey, 2016; Kim et al., 2016, Kristoufek, 2013; Lahmiri, 2011; Ron and Shamir, 2013) as new cryptocurrency journals are introduced and technical reports of central banks all over the world start to pay attention to this still relatively young phenomenon.

Acknowledgement

This research was supported by JSPS Grants-in-Aid Nos. 2538404, 2628089.

References

Akansu, A. N., Kulkarni, S. R. & Malioutov, D. M. (Eds.) (2016), Financial Signal Processing and Machine Learning, Wiley-IEEE Press.

Böhme, R., Christin, N., Edelman, B., & Moore, T. (2015) Bitcoin: Economics, Technology, and Governance, Journal of Economic Perspectives vol. 29, No. 2, pp. 213–238.

CoinDesk (2017) Bitcoin – Market Capitalization, http://www.coindesk.com/data/bitcoin-market-capitalization/, Accessed 2017/03/03.

Cortes, C. Vapnik, V. (1995) Support-vector networks, Machine Learning vol. 20, No. 3, pp. 273–297.

Elman, J. L. (1990) Finding Structure in Time, Cognitive Science vol. 14 No. 2, pp. 179–211.

Extance, A. (2015) Bitcoin and beyond, Nature vol. 526, Issue 7571, pp. 21-23.

Franzke, C. (2012) Predictability of extreme events in a nonlinear stochastic-dynamical model, Physical Review E vol. 85, article ID 031134.

Gyorfi, L., Ottucsak G. & Walk, H. (Eds.) (2012), Machine Learning For Financial Engineering (Advances in Computer Science and Engineering: Texts), Imperial College Press.

Hallerberg, S., Bröcker, J. & Kantz, H. (2008), Prediction of Extreme Events, in Nonlinear Time Series Analysis in the Geosciences, Lecture Notes in Earth Sciences vol. 112, pp. 35-59.

Houey, N. (2016) The Bitcoin mining game, Ledger, The Journal of Cryptocurrency and Blockchain Technology Research, vol. 1, pp. 53-68.

Kim Y. B., Kim J. G,. Kim W., Im J. H et al. (2016) Predicting Fluctuations in Cryptocurrency Transactions Based on User Comments and Replies. PLoS ONE vol. 11, No. 8, e0161197.

Kirikos, D. G. (2000) Forecasting exchange rates out of sample: random walk vs Markov Switching Regimes. Applied Economics Letters vol. 7, pp. 133-136.

Kristoufek L. (2013) BitCoin meets Google Trends and Wikipedia: Quantifying the relationship between phenomena of the Internet era. Scientific reports vol. 3, 3415.

Lahmiri, S. (2011) A comparison of PNN and SVM for stock market trend prediction using economic and technical information. International Journal of Computer Applications vol. 29, No. 3, pp. 24–30.

Ron D, Shamir A. (2013) Quantitative analysis of the full bitcoin transaction graph. Financial Cryptography and Data Security: Springer, pp. 6–24.