Microstructure and Market Dynamics in Crypto Markets

David Easley, Maureen O'Hara, Songshan Yang , and Zhibai Zhang*

April, 2024

We investigate the role of market microstructure metrics in explaining and predicting price dynamics for 5 major cryptocurrencies. Using machine learning, we show how microstructure measures of liquidity and price discovery have predictive power for price dynamics of interest for electronic market making, dynamic hedging strategies and volatility estimation. We identify important own market and cross-market effects for BTC and ETH Roll measures and VPINs. Our results are little changed during crypto winter, demonstrating a stability to these effects. Our findings suggest that market dynamics of cryptocurrencies can be viewed as similar to those of other investible asset classes.

*David Easley, Departments of Economics and Information Science, Cornell University; Maureen O'Hara, Johnson College of Business, Cornell University; Songshan Yang, Center for Applied Statistics and Institute of Statistics and Big Data, Renmin University of China; Zhibai Zhang, the Tandon School of Engineering, New York University

Microstructure and Market Dynamics in Crypto Markets

The challenge confronting cryptocurrencies has long been apparent: How to make crypto an asset class investible by retail and institutional traders? The resurgence of crypto prices suggests that some impediments to reaching this goal are receding, at least as they apply to retail investors. Following the January 2024 SEC approval of spot-market based Bitcoin ETFs, inflows to the new Bitcoin ETFs reached almost $70 billion in just 2 months.[1] This remarkable growth testifies to the appeal that being able to buy bitcoin exposure through brokerage accounts rather than via crypto exchanges or futures markets has for retail traders. But other obstacles remain for both retail and institutional investors.[2] Whether these difficulties can be overcome remains to be seen, but certainly fundamental to greater participation is the ability to understand what drives the market dynamics of cryptocurrencies.

Understanding these market dynamics, however, is not straightforward. As we discuss below, there is a developing literature attempting to understand crypto valuation drawing on standard tools from other asset classes. Other approaches focus on crypto-specific factors such as underlying users or mining costs to explain crypto valuations. Yet other researchers eschew understanding the determinants of crypto value and instead turn to more "black box" machine learning to simply predict future prices. While these approaches yield interesting insights, accurate crypto valuation continues to prove elusive, underscoring the difficulty of making crypto more

---

[1] AUM numbers as of March 28, 2024 from https://cryptonews.com/news/assets-invested-in-crypto-etfs-and-etps-rise-359-in 2024.htm#:~:text=ETFGI%2C%20an%20independent%20research%20firm%2C%20reports%20assets%20invested, increase%20from%20%2415.12%20billion%20at%20end%20of%202023. See also "Bitcoin ETFs on Better Win Streak than 95% of Traditional Funds," Forbes, March 25, 2024.

[2] Vanguard, for example, will not offer crypto-products arguing bitcoin is "a speculation [rather] than an investment". JP Morgan and Goldman Sachs have expressed similar views. Institutional roadblocks include regulatory issues (e.g. is crypto a security?, know your customer compliance), market manipulation concerns, valuation issues, market fragmentation, excessive volatility, and lack of liquidity. For discussion, see blog.ionixxtech.com/top-5-challenges-institutional-investors-face-in-crypto-trading/.

accessible to a wider investing audience who, understandably, want to know whether the crypto asset they are considering buying is over- or under-valued.

In this paper, we offer a different approach to investigating crypto market behavior. Our particular focus is on a basic question: Can standard microstructure measures predict crypto market dynamics? Unlike valuation-based analyses, our approach is based more on understanding the liquidity and price discovery process involved in crypto trading. Microstructure theory provides various measures related to liquidity (for example, the Amihud measure), asymmetric information and toxicity (Kyle lambda, VPIN), and overall spreads and auto-correlations (Roll measure and Roll Impact measure) that have been shown to matter for liquidity and price dynamics in other asset markets. We estimate these variables for 5 major cryptocurrencies, giving insight into how these metrics differ from those found in more standard market settings.

We then use these measures to predict 5 outcomes of market price dynamics of particular interest for electronic market making, dynamic hedging strategies and volatility estimation. These outcome variables are the signs of the change in realized volatility, the change in auto-correlation of realized returns, the change in skewness of realized returns, the change in kurtosis of realized returns, and the change in the Jarque-Bera statistic. We follow the approach taken in Easley, Lopez de Prado, O'Hara, and Zhang (ELOZ) (2021) of using machine learning to ask if our microstructure features can predict our outcome labels in cryptocurrency markets.

Some readers might find it odd to care about crypto market liquidity and price dynamics rather than crypto valuation *per se*. Two reasons to do so are paramount. First, as noted in ELOZ, in high frequency markets how the market is structured turns out to be critical in predicting where the market is going. The less "efficient" the market, the more predictable it is, so understanding the efficiency of crypto trading matters. The second reason is that institutional and high frequency traders rely on algorithmic trading approaches to optimize trading strategies. These strategies, in turn, rely on predicting market dynamics to determine the optimal path for executing trades. Algorithmic trading is widely used in crypto markets, where trade bots allow both crypto "whales" and retail traders alike the ability to trade dynamically.[3] But the algorithmic strategies employed there generally rely on common market indicators, and not on the underlying microstructure variables that may drive the more sophisticated trade execution strategies found in more standard asset markets.[4] The labels we focus on are inputs to those quantitative strategies, so understanding their predictability is fundamental to attracting such traders to the crypto space. Perhaps the simplest way to characterize our interest is that we are asking: Are crypto markets really different, and if so, how?[5]

Our research design uses high frequency data from Binance (the largest crypto exchange) for cryptocurrencies of five major blockchain platforms: Bitcoin (BTC), Ethereum (ETH), Ripple (XRP), Solano (SOL), and Cardano (ADA)[6]. We are interested in both the effects of own microstructure variables for prediction (e.g. does the ADA Roll metric predict the sign of changes in future ADA volatility) and the cross effects for prediction (e.g. does Bitcoin VPIN predict the sign of changes in future Ethereum autocorrelation). Our interest in these cross-asset effects reflects the reality that high frequency trading behavior often involves complex multi-asset

[3] For a discussion of trade bots see Coinbase, "How trade bots work? where it is noted that "some common parameters bots use include price, time frame, and order volume, while common market indicators include moving averages (MAs), relative strength index (RSI), and more". Available at https://www.coindesk.com/learn/what-are-crypto-trading-bots-and-how-do-they-work/

[4] There is some discussion in the crypto industry of the potential use of microstructure measures, but no systematic evaluation of microstructure measures and their ability to predict characteristics of the distribution of returns. See for example, https://medium.com/@kryptonlabs/vpin-the-coolest-market-metric-youve-never-heard-of-e7b3d6cbacf1

[5] A related approach is taken by Kogan, Makarov, Niessner, and Schoar [2024] who investigate whether retail trading in crypto differs from retail trading in stocks and gold. These authors find substantial differences in retail crypto trading.

[6] Each blockchain has a native token in which transactions are charged. In addition, various non-native tokens can be deployed on a blockchain. Here we focus the study only on the native tokens, which are represented by the tickers in parentheses.

strategies. We use random forest machine learning to ascertain overall predictability and to determine which features (i.e. microstructure variables) are most important for understanding market dynamics. Our sample period spans January 2021 – July 2023, allowing us to investigate whether the predictability of crypto markets behaved differently during the "crypto winter" period.

Our research provides a number of results, three of which we highlight here:

First, we find surprisingly high values for the Roll Measure and VPIN in crypto markets relative to more standard equity and futures market settings. The greater serial correlation in crypto prices is consistent with more momentum-based trading, while the higher levels of VPIN are indicative of greater trade toxicity arising potentially from more asymmetric information.

Second, we find strong predictability of microstructure measures for future market price dynamics. Averaging across all currencies and variables, we find an AUC > .55, a very strong result consistent with deviations from market efficiency.[7] Focusing on the individual features, we find similar predictability (AUC ranging from .54 - .61) with the exception of skewness where we find no predictability (AUC=.50). Overall, we find that predictability is driven by the Roll measure and by VPIN. Own measures of Roll and VPIN matter for most of our predicted features, with other microstructure features having little or no importance for own market prediction. Cross effects are also important: BTC Roll and VPIN and ETH Roll and VPIN have strong predictive power for price dynamics across the other crypto currencies.

Third, a particular concern for institutional trading methods is the stability of these market dynamics. We use the natural experiment of "crypto winter" to investigate how these market dynamics are affected by different market regimes. Somewhat surprisingly, we find no effects of

[7]A standard metric in random forest algorithms to measure the performance of classification models is ROC, or the receiver operating classification curve. The AUC stands for "area under the ROC curve" and essentially captures the predictive ability of the machine learning algorithm. For more discussion of evaluation techniques for financial machine learning see Lopez de Prado (2018).

crypto winter on our results, suggesting a remarkable stability to the price dynamic effects uncovered here.

Overall, we show that the price dynamics of the crypto market respond to its market microstructure in ways similar to that of other financial assets. That the market exhibits inefficiency is not too surprising – crypto markets are relatively young, exhibit large volatility relative to other asset classes and are much less regulated, all features that undermine achieving price efficiency.[8] But even in mature markets, there can be predictability from microstructure variables. As ELOZ [2021] showed, futures markets also exhibit inefficiency, although not to the degree found here in the crypto markets. What is true in both market settings is that over short horizons liquidity is predictable, and thus potentially exploitable for optimizing trading strategies. Moreover, our results on the strong cross effects of the Roll measure, a metric capturing autocorrelation properties, and VPIN, a measure of market toxicity, convincingly support that there are common factors driving price dynamics across cryptocurrencies.[9]

Our research relates to several streams of research directed to understanding the market drivers of cryptocurrencies. A large literature in computer science uses machine learning to predict future crypto prices and returns (see, for example, Koker and Koutmos (2020); Jaquart, Kopke, and Weinhardt (2022); Cortese, Kolm, and Lindstrom (2023); Filippou, Rapach, and Thimsen (2024)). Our analysis also relies on machine learning but differs from the black box prediction methods typically used in that we are testing for the effects of particular model-based metrics on market dynamics. In common with these papers, we find strong predictability, and thus inefficiency, in crypto markets.

[8] See also Nimalendram et al [2021] who use variance tests to show cryptocurrency inefficiency and examine its relationship to market regulation.
[9] For a discussion of factor pricing in crypto see Cong, Karolyi, Tang and Zhao (2022).

A large finance literature focuses on the valuation of cryptocurrencies (see Pagnotta and Buraschi (2018), Cong, et al (2020), Bias et al (2022) for theoretical analyses). Much of the empirical literature draws on applying valuation methodologies to estimate cryptocurrency values. Erb (2020) investigates linkages between gold valuation and bitcoin valuation but concludes that "neither gold nor bitcoin are obvious inflation hedges, stores of value, or safe havens" and so are hard to value.[10] Liu and Tsyvinski (2018) argue that crypto has no exposure to common equity market or macro factors or to currency or commodity markets.[11] Liu, Tsyvinski and Wu [2021] develop a three-factor model based solely on crypto market, size and momentum to explain returns. Cong, Karolyi, Tang and Zhao (2021) augment these factors with value and network factors, yielding a five-factor model of crypto returns. Our analysis suggests that inter-dependencies across crypto currencies arising from information and liquidity dynamics may be potential factors to consider in future research.

Another substantial body of research focuses on the trading of crypto assets (for a general survey see Hallaburda et al (2022)). Early papers here include Easley, O'Hara, and Basu (2018), Huberman and Lasko (2021), and Cong He, and Li (2021), who analyzed the equilibrium role of transaction fees and miners in Bitcoin. Recent work analyzes trading across crypto exchanges, with Makorov and Schoar (2019; 2020) finding extensive arbitrage opportunities and Crepelliere et al (2022) documenting a decreasing trend in such opportunities. Our results are consistent with the market inefficiency found there, with our results on its persistence during crypto winter underscoring the importance of understanding the cryptocurrency microstructure. Also relevant is Makorov and Schoar (2022) who present intriguing evidence on the trading and network structure

[10] See Erb (2020) pg. 15.

[11] Related research investigates the linkages of crypto markets to other markets, with Iyer (2022) establishing increased interdependence between crypto and equity markets.

of the bitcoin blockchain. These authors find that the bitcoin ecosystem is still dominated by large and concentrated players, raising anew the issue of whether crypto currencies can appeal to a wider investment audience. Research directly addressing this issue includes Hardle, Harvey and Ruelle (2020), Harvey et al (2022) and Ang, Morris, and Savi (2022).

The paper is organized as follows. Section II describes the microstructure variables we consider, the market statistics we predict using these microstructure variables and the random forest procedure we use to generate predictions. Section III describes the data set: Binance data on prices and volumes of trade for the leading five crypto currencies for the period January 2021 to July 2023. Section IV provides results about predictability and the relative importance of each microstructure measure in generating predictions. Section V offers two robustness tests: the effect of crypto winter; and using logistic regressions rather than random forests to generate predictions. Section VI concludes.

II. Research Design

We ask whether five standard microstructure variables are useful in predicting various measures of market dynamics. The specific microstructure measures considered are: Roll measure, Roll impact measure, Kyle's lambda, Amihud measure and VPIN (Volume synchronized measure of information-based trade). These microstructure measures are intended to measure illiquidity or the presence of information-based trade. Illiquidity and information should lead to price volatility, and ELOZ (2021) demonstrate that these microstructure measures are successful in predicting price dynamics in futures markets.

The price dynamics we focus on are measures of changes in the distribution of realized returns. The specific price dynamics measures we predict are: the sign of the change in the sequential correlation, the sign of the change in the Jarque-Bera statistic, the sign of the change in realized volatility, the sign of the change in kurtosis, and the sign of the change in the skewness. These measures have distinct implications for trading strategies. For example, if realized volatility is expected to increase, then increasing the speed of algo execution would be expected to reduce fill price uncertainty. An increase in predicted serial correlation can result in greater or less price impact depending upon the trade side being executed. As ELO [2015] show, this should change the optimal speed of trading.[12] The Jarque-Bera statistic captures normality of returns, so its increase signals non-normal returns, suggesting that estimates of implementation shortfall may be too small. If skewness is expected to increase, then the distribution of returns is shifted to one side, perhaps consistent with toxicity in order flow. An increase in kurtosis means greater weight in the tails, an outcome that may signal a withdraw of liquidity support by market makers. Delaying the speed of execution would then be optimal.

These price variables are defined using Binance data on prices and volume. For each crypto we first create one-minute time bars; that is, we split the data into segments of length one-minute and we record the price at the beginning of each time bar, at the end of each time bar, and the dollar volume of trade that occurs within each time bar.[13] Let t=1,2,… index time measured in minutes. The basic variables we consider are the ending price p_t in time bar t, the return $r_t = (p_t - p_{t-1}) / p_{t-1}$ in time bar t and V_t the dollar volume of trade in time bar t.

We next compute a realization of each of our microstructure measures for each time bar t. Each of these microstructure measures are based on some amount of past data. For any time bar t, a microstructure measure at t is computed using data in periods t, t-1, …, t-W where W is the

[12] These authors also demonstrate how incorporating microstructure variables into trading strategies can improve upon the outcomes provided by standard trade algorithms. In particular, an algorithm based on predicted VPIN changes and volume participation dominates a VWAP trading strategy. See Lopez de Prado et al (2020) for more discussion.

[13] We use one-minute time bars because that is the highest granularity bar data the public exchanges offer. This time period also seems appropriate for capturing market dynamics in a high frequency setting.

lookback window. In our analysis we consider lookback windows of 50 bars and 100 bars. For example, the Amihud measure at time bar t is the average ratio of absolute returns to dollar volume where the average is computed over the past W time bars. Note that a microstructure measure at time bar t and one at time bar t+1 are computed using W-1 overlapping bars of market data.

Our market microstructure variables are defined from these basic variables as follows:

1. The Roll measure is

$$2\sqrt{|cov(\boldsymbol{\Delta P_t}, \boldsymbol{\Delta P_{t-1}})|},$$

$$\boldsymbol{\Delta P_t} = [\Delta p_{t-W}, \Delta p_{t-W+1}, \dots, \Delta p_t],$$

$$\boldsymbol{\Delta P_{t-1}} = [\Delta p_{t-W-1}, \Delta p_{t-W}, \dots, \Delta p_{t-1}],$$

where $\Delta p_t = p_t - p_{t-1}$.

2. The Roll impact measure---the Roll measure divided by dollar volume over a certain period, is

$$2 \frac{\sqrt{|cov(\boldsymbol{\Delta P_t}, \boldsymbol{\Delta P_{t-1}})|}}{p_t V_t}.$$

3. The Amihud measure is

$$\frac{1}{W}\sum_{i=t-W+1}^{t}\frac{|r_i|}{p_i V_i},$$

4. Kyle's λ is

$$\frac{p_t - p_{t-W}}{\sum_{i=t-W}^{t} b_i V_i},$$

where $b_i = sign(p_i - p_{i-1})$.

5. VPIN is

$$\frac{1}{W}\sum_{i=t-W+1}^{t}\frac{|V_i^S - V_i^B|}{V_i},$$

where $V_i^B = V_i Z\left(\frac{\Delta p_i}{\sigma_{\Delta p_i}}\right), V_i^S = V_i - V_i^B$.

We use these microstructure measures to predict the signs of changes in various market statistics. These "signs of change" are represented by -1 for a negative change and +1 for a positive change. Thus, our market statistics data is a sequence of -1 and +1, one for each time bar. Similar to the way we compute microstructure measures, these market statistics are also computed using some number of past observations. For example, the distribution of realized returns as of time bar t is the empirical distribution of returns over some number of past time bars. The number of past time bars used in computing each market statistic is also W (the lookback window). Note that this procedure implies that the sign of change in a market statistic at time bar t and at time bar t+1 also uses W-1 overlapping observations of data. So, when we want to predict the sign of the change in a market statistic we predict it over enough future time bars to avoid using overlapping data. That is, we use a substantial ``look ahead'' window. We set this look ahead window to be h=1,500 bars; roughly one day of trading.

Formally the signs of change in market statistics are:

1. The sign of the change in realized volatility

$$sign(\sigma_{t+h} - \sigma_t)\ ,$$

   where σ_t is the realized volatility of one bar returns over a look back window of size W.

2. The sign of change in Jarque-Bera statistics of realized returns

$$sign\big(JB(r_{t+h}) - JB(r_t)\big)\ ,$$

$$JB(r) = \frac{n}{6}\left(S^2 + \frac{1}{4}(C-3)^2\right),$$

where S is the skewness and C is the kurtosis of realized returns r over the previous W bars.

3. The sign of change in sequential correlation of realized returns

$$sign(sc_{t+h} - sc_t),$$

$$sc_t = corr(r_t, r_{t-1}).$$

4. The sign of change in absolute skewness of realized returns:

$$\text{sign}[Skew_{t+h} - Skew_t].$$

5. The sign of the change in kurtosis of realized returns:

$$\text{sign}[Kurt_{t+h} - Kurt_t].$$

The microstructure literature suggests that a financial asset's microstructure measures should provide information about its price dynamics. But the literature provides little guidance about exactly what form this relationship should take. From prior research it also seems reasonable to expect cross-asset relationships, for example BTC or ETH microstructure measures may help predict market statistics of other cryptos, but again the literature provides no guidance about what form this relationship should take.[14] To avoid restricting ourselves to an arbitrary structure we use machine learning to discover these relationships.

A wide variety of machine learning procedures could be used to ask about the predictive content of our market microstructure measures (see Lopez de Prado (2018) and Kelly and Xiu (2023) for discussion of these approaches). We chose a random forest procedure primarily to make it possible to compare our results to previous results about the predictive content of microstructure measures. In ELOZ (2021) we used a random forest applied to futures data to ask if these same

[14] See, for example, research on crypto currency interdependence by Kukacka and Kristofex (2023) and Qureshi et al (2020).

microstructure measures could predict changes in market statistics. In this standard financial market, the answer was yes; some own market measures matter and so do some cross-market measures. Here we ask if these same microstructure measures provide insight into crypto price dynamics.

In the language of random forests, our "features" (the microstructure measures) are used to predict "labels" (the market statistics). Our random forest procedure begins by building, for each crypto currency, decision trees based on repeated cuts of the time series of labels and features into two pieces. The division of the data is determined by a two-step process. First, for each feature we compute the information gain (gain in homogeneity of the data) that would be obtained by splitting the data using that feature. Second, we cut the data into two pieces (two branches of the decision tree) by randomly selecting two features and using the feature with the largest information gain. We repeat this cutting process until no additional cut will yield an information gain. For each crypto currency this process yields a decision tree that predicts labels for any assignment of features. To avoid tying our predictions to a single tree that may be too heavily influenced by randomness in the data, we create multiple decision trees for each crypto currency. This is done by creating 100 trees using bootstrapped samples of the data from the actual data set. The prediction given any list of features is then the majority prediction using the collection of trees.[15]

Ultimately, we are interested in the accuracy of our predictions and in which features contribute to any accuracy.[16] In a regression analysis it is standard to compute p-values to measure the importance of any predictor. This is, however, an in-sample measure and its computation is tied to the regression analysis. In contrast, the feature importance measure Mean Decreased

[15] For more discussion of this procedure see ELOZ (2021).
[16] For a discussion of how best to measure accuracy in financial machine learning applications see Lopez de Prado (2018), Chapter 8.

Accuracy (MDA) is based on out-of-sample prediction, and it can be used to compute the importance of any predictive variable. Mean Decreased Accuracy for a feature represents how much prediction accuracy we lose if we compute accuracy first using that feature and then using shuffled values of the feature. To compute MDA, we first split the data into disjoint training and test sets. We run our random forest procedure on the training set and compute accuracy of predictions on the test set (the definition of accuracy is given below). We then rerun the random forest procedure on the training set after randomizing the value of a feature and compute accuracy again. The relative loss in accuracy is defined as the MDA measure for this feature.

III. Data

The data used in this study are obtained from the Binance public database for the period January 2021 to July 2023. Historically, there have been a large number of crypto exchanges, some of which were relatively short-lived. As of December 2023, active centralized crypto exchanges with daily volume greater than one-hundred million dollars numbers in the dozens, according to the crypto analytics site https://www.coingecko.com/. Among these exchanges, Binance is the largest as measured by daily dollar volume so we chose it for the focus our study. Because each exchange runs their own limit order book, we cannot directly apply our current analysis to data combined from multiple exchanges.

The format of the data is the standard time-bar candle-stick with 1-minute intervals, including open, close, high, low price and volume aggregated from tick data within each interval. All features and labels in this study are derived from this data. Since the number of cryptocurrencies has grown exponentially, and is quite volatile, it is infeasible to examine all of the tokens traded on the exchange. We focus on the top 5 cryptocurrencies as measured by their market-capitalization in January 2021: BTC, ETH, ADA, SOL, and XRP. A deep dive into how and if trade in these dominant tokens differs from the smaller market size tokens would be an intriguing topic for future research.

IV. Results

Table 1 provides estimates of the mean values of the market microstructure variables over our sample period. As is standard in crypto settings, we express each currencies' microstructure measures with respect to the U.S. dollar as given by Tethers (USDT). As expected, all variables have positive signs. Averaging over longer time bars tends to reduce mean values for the Roll and Kyle metrics, but has little impact on the Amihud and VPIN measures. The Roll measure is markedly different across the currencies, with Bitcoin having substantially higher autocorrelation. The presence of "whales" trading large quantities algorithmically could be one reason for such a finding. The VPIN measure is remarkably stable across currencies, but its level is surprising: ELO (2012) found (using a slightly different methodology) average VPINs for the E-mini S&P500 and crude oil futures of 0.22 to 0.23 whereas these are range from .45 to .47.[17] Such high toxicity is consistent with greater information-based trading in crypto markets.

We now ask how much predictive accuracy we can achieve from our microstructure measures and which microstructure measures contribute to this predictive accuracy. We consider two ways to measure predictive accuracy. One measure (Accuracy) is the number of correct predictions of the sign of change divided by the total number of predictions. The second measure (Area Under the Curve or AUC) can be interpreted as the probability that our fitted random forest will rank a randomly drawn +1 higher than it ranks a randomly drawn -1. We confine our discussion in the text to the simple Accuracy measure as the two measures yield nearly identical

[17] They are not as high as the highest VPIN levels of approximately 0.8 that ELO (2012) found for the E-mini during the time period of the flash crash.

results.[18] We focus on the entire sample period (January 2021 through July 2023); in a later section we divide the sample into two periods and ask about the effects of Crypto Winter.

For each crypto currency we use the random forest procedure described in Section II applied to 25 features: the five microstructure measures for each of the five crypto currencies. So, we include both the effect of own microstructure measures and cross-currency microstructure measures. Random guessing should lead to predictive accuracy of approximately 0.5 as we are predicting whether labels are positive (+1) or negative (-1). Anything above 0.5 suggests that our features have some power in predicting our labels. We do not attempt to maximize prediction accuracy. For our purposes, it is enough to demonstrate that our microstructure measures contribute to substantial prediction accuracy here as they do in more established financial markets. Previous research (ELOZ 2020) on futures markets using similar features to predict similar labels found average predictive accuracy in the range 0.54 to 0.61.

We first examine average (averaged over the five crypto currencies) predictive accuracy. The most important question is whether we can make useful predictions at all. A secondary question is how the level of predictive accuracy compares to that found in other financial markets. Note that we are predicting changes in market statistics one day ahead and we make predictions every minute. So even small amounts of accuracy (above 0.5) can be valuable. Average (across all cryptos in our sample and all labels) predictive accuracy is provided in Table 2. Depending on the number of bars used as a lookback window in constructing our market microstructure measures, we find that predictive accuracy is between 0.53 and 0.54.[19]

If we exclude skewness, predictive accuracy for our entire sample period ranges (across accuracy measures, the number of bars used and labels) from 0.52 to 0.58. These results are

[18] Results for both accuracy measures are provided in the tables.
[19] The results for AUC are also given in Table 1 and they range from 0.53 to 0.54.

provided in the last two rows of the two final columns of Tables 3 (Auto-correlation), 4 (JB statistic), 5(Kurtosis), and 6 (Realized Volatility). However, as the last two columns of Table 7 show, skewness is not predictable; for skewness we find predictive accuracy of 0.5 meaning that for the crypto currencies we examine our microstructure measures are not at all useful in predicting market statistics for skewness. This is similar to results for futures where skewness is the least predictable label.

Predictive accuracy is greatest for the sign of the change in realized volatility. For both of accuracy measures, the average accuracy of prediction ranges from 0.56 to 0.58 depending on the number of bars used. This is a remarkably high accuracy level for random forest predictions in financial applications, see Lopez de Prado (2018).

The predictive accuracy results indicate that the crypto market exhibits inefficiencies. When order flow is imbalanced, causing positive correlation in price changes and an increased Roll measure, volatility increases and it remains high through our look-ahead window of 1,500 bars or roughly a day. The ability to predict changes in realized volatility from microstructure measures intended to measure imbalances or correlations in order flow suggests that there is trend following in these markets which could be exploited by sophisticated trading algorithms.

We find it interesting that the level of predictive accuracy is not very different from that found in futures markets using a similar approach to prediction. Both market settings are electronic and trade almost continuously over a 24-hour day[20], but they do differ in that futures are well-established markets widely used by institutional traders whereas crypto trading is more nascent and dominated by crypto "natives". Our results suggest that the market dynamics of crypto

[20] CME, one of the leading futures markets offer products that trade up to 23 hours a day Monday through Friday. Many crypto exchanges operate 24 hours, 7 days a week without a break.

markets may already be similar enough to more mature markets to facilitate analogous sophisticated trading techniques.

Equally important for our purposes is determining which microstructure measures contribute to prediction accuracy. Figures 1 through 4 provide average MDA scores (averaged over our five crypto currencies) for prediction of the change in realized volatility (Figure 1), the change in Kurtosis (Figure 2), the change of the JB statistic (Figure 3) and the change in auto-correlation (Figure 4).[21] These results are derived from prediction of each market statistic for each crypto currency using only that currencies own microstructure measures. These results are remarkably consistent across the various market statistics. For prediction of each market statistic, the Roll measure is the most important feature as measured by MDA. VPIN is the second most important measure and the Roll impact measure is third most important. It is reasonable to expect that microstructure measures for other crypto currencies could be useful in predicting market statistics for a specific crypto. In particular, as BTC and ETH are the leading cryptos it could be that trade in these cryptos leads trade, and thus changes in market statistics, for other cryptos. Figures 5 through 9 provide the results for this analysis. Each figure provides MDA scores for one of our price dynamics labels for each of our five crypto currencies. For example, Panel 1 of Figure 5 provides MDA scores for the 25 microstructure measures for the sign of the change in realized volatility for ADA. Examining this panel of Figure 5 shows that ADA's own Roll measure is the most important feature. The next three most important features for ADA are, in order, ADA's own VPIN measure ,the Roll measure for BTC, and the Roll measure for ETH. and

Figures 5 through 9 tell a compelling story. In almost every case, each crypto's own Roll measure is the most important feature for predicting price dynamics. It is reassuring that some

[21] We do not provide results for Skewness as there is no predictive accuracy for it.

microstructure measure matters, but the details of prediction in crypto markets are different from those in futures markets. In futures markets, the Amihud measure and VPIN were the most important own measures. In our crypto sample, own VPIN shows up frequently as an important feature; and only occasionally do the own Roll impact measure or the Amihud measure have importance. The own Kyle measure typically has a low, nearly-zero MDA score.

The cross-market result for Bitcoin and Ethereum reveal particularly interesting dynamics. The BTC Roll measure and BTC VPIN are the most important features in driving Bitcoin price dynamics, with the Ethereum Roll measure also playing an important but lesser role. However, the other crypto currency cross-measures have virtually no influence. Similarly, price dynamics for Ethereum are driven by the ETH Roll measure and the ETH VPIN with the Bitcoin Roll measure playing again playing an important but lesser role. As is the case for Bitcoin, the other crypto currency cross-measures have virtually no influence. Across all other cryptos and labels we study, Roll measures for BTC and ETH have strong predictability signified by high MDA scores. All other cross-crypto features typically have very low MDA scores. This suggests that trade in BTC and ETH leads price changes and volatility in other cryptos; not a surprising result, but one that is consistent with the notion that trade in other cryptos follows trade in these two large cryptos.

Although the level of predictability we find suggests that there are inefficiencies in the crypto markets, the fact that market microstructure measures are important for price dynamics demonstrates that these markets have much in common with more standard financial asset markets. Our random forest analysis does not imply causation flowing from our features to our labels. But it does suggest that trading tools based on own Roll and VPIN as well as cross BTC and ETH Roll measures can be valuable.

V. Crypto Winter and Other Robustness Tests

In this section we consider two robustness tests. First, we ask if our results are stable over time using the beginning of crypto winter as a date to split our sample. Second, we ask if a logistic regression would produce similar results.

V.1 . Crypto Winter

Prices and trading volumes of crypto currencies changed dramatically over our sample period. Approximately the first half of our sample falls in a “boom” period for crypto currencies with rising prices and increasing high trading volumes. During this period, overall daily trading volume reached a peak of $158.64B on April 10, 2021, and remained high over much of 2021. Emblematic of crypto prices, during this period the price of Bitcoin reached a high of $67,617 on November 2021, where after prices began a steady decline, reaching a low of $15,742 in October 2022.[22] This latter period is generally referred to as “crypto winter”. A natural question is whether our predictability results differ in the period before crypto winter and after it began. We selected November 10, 2021 as the date to use in breaking our sample into two pieces as the total crypto market capitalization reached its high on November 9, 2021.

We reran our analysis separately on these before and after periods. *Ex ante* it is not obvious whether predictability should increase or decrease as a result of crypto winter or if there should be changes which microstructures contribute most to them. Crypto winter was a time of great uncertainty in the value of crypto and this should make predicting more difficult. But the market was also likely less efficient and this might make predicting easier.

Table 8 provides our results. Although valuations changed dramatically between these two sub-periods, predictability is nearly unchanged and the importance of our various microstructure measures is also unchanged. These results suggest that although valuations changed, the structure

[22] Along with depressed prices, crypto winter also featured several spectacular failures of crypto currencies and exchanges. For discussion of these aspects of crypto winter see Arner, Zetsche, Buckley, and Kirkwood (2023).

of trading did not change, and the level of inefficiency in the market was also nearly unchanged. We regard this as good news for Crypto markets as it implies that trading tools based on these microstructure measures should be robust to the extreme volatility of Crypto currency markets.

Furthermore, in Tables 9 and 10 we present the values calculated for before and after November 2021 for the microstructure features used in our analysis. It is worth noting that out of the five market microstructure variables that we consider, Roll measure, Amihud measure and Kyle's λ are proportional to the scale of price while Roll Impact and VPIN are not. As such, because the price of most crypto tokens dropped significantly after the peak in November 2021, there is a more pronounced change in the Roll measure, the Amihud measure and Kyle's λ. On the other hand, the values of Roll Impact and VPIN are comparably more stable before and after November 2021.

V.2. Logistic Regression

An alternative to using a random forest to predict our binary labels is to use a logistic regression. The logistic regression uses a liner regression to model the log odds of the two labels.[23] The aggregated accuracy and MDA results for the logistic regression are provided in Figure 10. Overall, the results are similar to those obtained via a random forest. There is predictive accuracy for all labels other than skewness. The ranking of features by MDA is unchanged, although the actual MDA scores for the most important features, the Roll measure and VPIN, are increased relative to the scores for the less useful features. This is reassuring as these results suggest that our ability to predict and our ranking of the importance of various microstructure measures is not a result of the specific method used to classify observations of market statistics.

[23] The specific approach we use is discussed in ELOZ (2022), Section 4.6.

VI. Conclusion

The degree of predictability for market dynamics in leading crypto currencies indicates that there are some inefficiencies in the crypto markets, and it's reasonable to suspect that these inefficiencies would be even greater for less well established cryptos and might differ across exchanges. However, at least for trades of these five leading cryptos in Binance, the amount of predictability is not particularly different from what we found in futures markets. And this predictability remained almost unchanged during the crypto winter period, suggesting that market dynamics on the largest crypto exchange exhibit substantial stationarity. Specifically, using standard microstructure metrics, we find non-trivial prediction accuracy and AUC scores for multiple return statistics measures such as volatility and auto-correlation. We also highlight the prominent predictive role of auto-correlation (captured by the Roll Measure), providing a microstructure foundation for the momentum observed in crypto prices.

Perhaps the most intriguing result is that the market microstructure measures we find to be important for price dynamics in cryptos are similar to those that matter for prediction in futures markets. This similarity suggests that these crypto markets have much in common with liquidity and price dynamics in more standard financial asset markets. Our random forest analysis does not imply causation flowing from our features to our labels. But it does suggest that trading tools based on own Roll and VPIN as well as cross BTC and ETH Roll measures can be valuable. For institutional and high frequency traders, whose trading relies on sophisticated algorithmic and optimized trading strategies, this commonality removes an important obstacle to their participation. Whether that tips the balance to crypto becoming an asset class in its own right remains to be seen.

References

Ang, Andrew and Morris, Tom and Savi, Raffaele, 2022, Asset Allocation with Crypto: Application of Preferences for Positive Skewness, Available at SSRN: https://ssrn.com/abstract=4042239 or http://dx.doi.org/10.2139/ssrn.4042239

Arner, D. W. , Zetzsche, D. Buckley, R. and J. Kirkwood, 2023, The Financialization of Crypto: Lessons from FTX and the Crypto Winter of 2022-2023, ,Available at SSRN: https://ssrn.com/abstract=4372516 or http://dx.doi.org/10.2139/ssrn.4372516

Biais, Bruno and Bisiere, Christophe and Bouvard, Matthieu and Casamatta, Catherine and Menkveld, Albert J., 2023, Equilibrium Bitcoin Pricing, Journal of Finance, 78(2) 967-1014

Cong, Lin and He, Zhiguo and Li, Jiasun, 2021, Decentralized Mining in Centralized Pools, Review of Financial Studies, 34(3)1191-1235.

Cong, Lin W. and Karolyi, George A. and Tang, Ke and Zhao, Weiyi, 2022, Value Premium, Network Adoption, and Factor Pricing of Crypto Assets, Available at SSRN: https://ssrn.com/abstract=3985631 or http://dx.doi.org/10.2139/ssrn.3985631

Cong, Lin and Li, Ye and Wang, Neng, 2020, Tokenomics: Dynamic Adoption and Valuation Review of Financial Studies, 34 (3), 1105-1155.

Cortese, Federico and Kolm, Petter N. and Lindstrom, Erik, 2023, What Drives Cryptocurrency Returns? A Sparse Statistical Jump Model Approach. Available at SSRN: https://ssrn.com/abstract=4330421 or http://dx.doi.org/10.2139/ssrn.4330421

Crépellière, Tommy and Pelster, Matthias and Zeisberger, Stefan, 2022, Arbitrage in the Market for Cryptocurrencies Journal of Financial Markets, forthcoming

Easley, D., M. Lopez de Prado, and M. O'Hara, ELO, 2012, Flow Toxicity and Volatility in a High Frequency World, *Review of Financial Studies*, 25, 5, 1457-93.

Easley, D., M. Lopez de Prado, and M. O'Hara, ELO, 2015, "Optimal Execution Horizon," Mathematical Finance, , 25(3), 640-672.

Easley, D., M. Lopez de Prado, M. O'Hara, and Z. Zhang ELOZ, 2021. Microstructure in the machine age. *Review of Financial Studies* 34:3316–63.

Easley, D., M. O'Hara, and S. Basu, 2019, From Mining to Markets: The Evolution of Bitcoin Transaction Fees, Journal of Financial Economics, 134(1), 91-109.

Erb, Claude B., 2020, Bitcoin is Exactly Like Gold Except When it Isn't, Available at SSRN: https://ssrn.com/abstract=3746997 or http://dx.doi.org/10.2139/ssrn.3746997

Filippou, Ilias and Rapach, David and Thimsen, Christoffer, 2024, Cryptocurrency Return Predictability: A Machine-Learning Analysis, Available at SSRN: https://ssrn.com/abstract=3914414 or http://dx.doi.org/10.2139/ssrn.3914414

Halaburda, Hanna and Haeringer, Guillaume and Gans, Joshua and Gandal, Neil, 2022 The Microeconomics of Cryptocurrencies, Journal of Economic Literature, 60(3) 971-1013.

Hardle, W., C. Harvey and R. Reule, 2020, Understanding Cryptocurrencies, Journal of Financial Econometrics, 12:2, 181-208.

Harvey, C., T.A. Zeid, T. Draaisma, M. Luk, H., Neville, A. Ryzm, and O. Van Hemert, An investor's guide to Crypto, 2022, Available at SSRN: https://ssrn.com/abstract=4124576 or http://dx.doi.org/10.2139/ssrn.4124576

Huberman, G., J. Leshno, and C. Moallemi, 2021, "Monopoly without a Monopolist: An Economic Analysis of the Bitcoin Payment System," Review of Economic Studies, 88 (6), 3011-3040;

Iver, T., 2022, Crypto Connections; Spillover between crypto and equity markets, IMF Global Stability Note No. 2022/1.

Jaquart, P., S. Kopke, and C. Weinhardt, 2022, Machine learning for cryptocurrency market prediction and trading, Journal of Finance and Data Science, 8, (2022) 331-352.

Kelly, Bryan T. and Xiu, Dacheng, 2023, Financial Machine Learning, Available at SSRN: https://ssrn.com/abstract=4501707

Kogan, S., I. Makarov, M. Niessner, and A. Schoar, 2024Are Cryptos Different? Evidence from Retail Trading, Journal of Financial Economics, forthcoming. Journal

Koker, T. and D. Koutmos, Cryptocurrency trading using machine learning, 2020, *J. Risk Financial Manag.* , *13*(8), 178

Kukacka, Jiri and Kristoufek, Ladislav, 2023, Fundamental and Speculative Components of the Cryptocurrency Pricing Dynamics, Financial Innovation (9).

Liu, Yukun, Tsyvinski, Aleh, and Wu, Xi, 2022, Common Risk Factors in Cryptocurrency, Journal of Finance, 77(2) 1133-1177.

Lopez de Prado, M., 2018, Advances in Financial Machine Learning, (Wiley; New York).

Makarov, Igor, and Antoinette Schoar. 2019. "Price Discovery in Cryptocurrency Markets." AEA Papers and Proceedings, 109: 97-99**.**

Makarov, Igor and Schoar, Antoinette, 2020, Trading and Arbitrage in Cryptocurrency Markets, Journal of Financial Economics, 135(2) 293-319.

Makarov, Igor and Schoar, Antoinette, 2020, Blockchain Analysis of the Bitcoin Market Working paper, Available at SSRN: https://ssrn.com/abstract=3942181 or http://dx.doi.org/10.2139/ssrn.3942181

Nimalendran, Mahendrarajah and Pathak, Praveen and Petryk, Mariia and Qiu, Liangfei, Informational Efficiency of Cryptocurrency Markets (February 11, 2021). Available at SSRN: https://ssrn.com/abstract=3818818

Pagnotta, E. and Buraschi, A. (2018). An equilibrium valuation of bitcoin and decentralized network assets. Working paper, Imperial College.

Qureshi, S., M. Aftab, E. Bouri, and T. Saeed, 2020, Dynamic interdependence of crypto markets: An analysis across time and frequency, PhysicsA: Statistical Mechanics and its Applications, 559(1).

Tables

Table 1: Mean values of market microstructure variables

Type	Window	Roll Measure	Roll Impact	Kyle's Lambda	Amihud	VPIN
SOLUSDT	50 bars	0.073091217	1.1765E-06	8.60909E-05	3.17524E-08	0.469121339
	100 bars	0.06355549	1.06911E-06	4.54613E-05	3.0913E-08	0.462114902
XRPUSDT	50 bars	0.000640782	5.43329E-09	2.73E-09	8.80316E-09	0.465487876
	100 bars	0.000547614	5.05243E-09	6.91752E-09	9.06127E-09	0.456565606
ADAUSDT	50 bars	0.001026723	2.5131E-08	6.55367E-09	3.04002E-08	0.464679
	100 bars	0.000923532	2.33606E-08	1.77393E-09	2.95753E-08	0.458025021
BTCUSDT	50 bars	20.36453232	1.78623E-05	15.9869968	4.28419E-10	0.469240396
	100 bars	18.16115194	1.62215E-05	1.255397095	4.28417E-10	0.459877142
ETHUSDT	50 bars	1.643139356	2.83363E-06	0.002306957	1.19623E-09	0.470527845
	100 bars	1.45823823	2.58103E-06	0.00490705	1.19626E-09	0.461889666

Table 2: Aggregated accuracy and AUC

Window	Aggregated Accuracy	Aggregated AUC
50 bars	0.538089	0.538134
100 bars	0.530436	0.530428

Both accuracy measures are aggregated across all 5 cryptocurrencies, across the entire test period, and across all 5 labels. The results provided in the table are the average values.

Table 3: Feature importance and prediction performance for auto-correlation

Period	Window	Roll Measure	Roll Impact	Kyle's Lambda	Amihud	VPIN	Accuracy	AUC
2021.1 – 2021.11	50 bars	0.04324	0.01002	0.00444	0.00364	0.01496	0.533329	0.533324
	100 bars	0.03972	0.00594	0.00452	0.00244	0.01042	0.527921	0.527826
2021.11 – 2023.7	50 bars	0.04506	0.01044	0.00168	0.00298	0.01674	0.536524	0.536573
	100 bars	0.03606	0.00722	0.00112	0.00164	0.01396	0.526781	0.526674
2021.1- 2023.7	50 bars	0.04426	0.00986	0.00396	0.00464	0.01552	0.531708	0.531698
	100 bars	0.03892	0.00630	0.00424	0.00318	0.01016	0.523608	0.523532

Table 4: Feature importance and prediction performance for JB Statistics

Period	Window	Roll Measure	Roll Impact	Kyle's Lambda	Amihud	VPIN	Accuracy	AUC
2021.1 – 2021.11	50 bars	0.04662	0.00838	0.00502	0.00006	0.0154	0.532057	0.532077
	100 bars	0.04256	0.00544	0.00516	-0.0002	0.00884	0.530776	0.530836
2021.11 – 2023.7	50 bars	0.04918	0.00772	0.00036	0.00142	0.01678	0.527283	0.527297
	100 bars	0.03978	0.00518	0.00106	0.00042	0.01404	0.523098	0.523367
2021.1- 2023.7	50 bars	0.04822	0.00814	0.00476	0.00124	0.01636	0.532994	0.533029
	100 bars	0.04330	0.00616	0.00506	0.00190	0.00946	0.531124	0.531142

Table 5: Feature importance and prediction performance for Kurtosis

Period	Window	Roll Measure	Roll Impact	Kyle's Lambda	Amihud	VPIN	Accuracy	AUC
2021.1 – 2021.11	50 bars	0.03308	0.00814	0.0036	0.00256	0.01882	0.544144	0.544135
	100 bars	0.03196	0.0048	0.00438	0.00198	0.01082	0.531292	0.531236
2021.11 – 2023.7	50 bars	0.0335	0.00852	0.00214	0.00344	0.0221	0.545655	0.545724
	100 bars	0.02728	0.00586	0.00184	0.00192	0.01498	0.529588	0.529195
2021.1- 2023.7	50 bars	0.03396	0.00792	0.00332	0.00328	0.01888	0.544184	0.544134
	100 bars	0.03124	0.00510	0.00452	0.00286	0.01066	0.530874	0.530842

Table 6: Feature importance and prediction performance for Realized Volatility

Period	Window	Roll Measure	Roll Impact	Kyle's Lambda	Amihud	VPIN	Accuracy	AUC
2021.1 – 2021.11	50 bars	0.05662	0.00922	0.00562	-0.00092	0.00808	0.579853	0.57897
	100 bars	0.04578	0.0048	0.0036	-0.00226	0.0057	0.562562	0.561153
2021.11 – 2023.7	50 bars	0.06442	0.00822	0.00092	0.00132	0.00922	0.584063	0.582128
	100 bars	0.05098	0.00542	0.00182	0.00014	0.00962	0.567168	0.564179
2021.1– 2023.7	50 bars	0.058120	0.00878	0.00534	0.001060	0.00922	0.581566	0.581797
	100 bars	0.046740	0.00562	0.00376	0.000660	0.00664	0.564701	0.564723

Table 7: Feature importance and prediction performance for Skewness

Period	Window	Roll Measure	Roll Impact	Kyle's Lambda	Amihud	VPIN	Accuracy	AUC
2021.1 – 2021.11	50 bars	0.03448	0.00804	0.00348	0.00266	0.01202	0.500303	0.50031
	100 bars	0.03202	0.00452	0.00366	0.00214	0.00858	0.501256	0.501286
2021.11 – 2023.7	50 bars	0.03704	0.00858	0.00162	0.00318	0.01402	0.504024	0.50409
	100 bars	0.02936	0.006	0.00138	0.00124	0.01174	0.503889	0.503985
2021.1– 2023.7	50 bars	0.035320	0.00782	0.00302	0.00342	0.01238	0.499993	0.500012
	100 bars	0.031440	0.0051	0.00392	0.00282	0.00862	0.501874	0.501902

Table 8: Aggregated accuracy and AUC before and after Nov 2021.

Period	Window	Aggregated Accuracy	Aggregated AUC
2021.1 – 2021.11	50 bars	0.537937	0.537763
	100 bars	0.530761	0.530467
2021.11 – 2023.7	50 bars	0.53951	0.539162
	100 bars	0.530105	0.52948
2021.1 – 2023.7	50 bars	0.538089	0.538134
	100 bars	0.530436	0.530428

Both accuracy measures are aggregated across all 5 cryptocurrencies, and across all 5 labels. The results in the table are the averages.

Table 9: Mean values of market microstructure variables before Nov 2021

Type	Window	Roll Measure	Roll Impact	Kyle's Lambda	Amihud	VPIN
SOLUSDT	50 bars	0.095517402	1.17672E-06	7.20968E-05	3.46395E-08	0.472742998
	100 bars	0.095411928	1.05633E-06	2.02224E-05	3.12291E-08	0.467878084
XRPUSDT	50 bars	0.001119395	3.97082E-09	1.98885E-09	4.44686E-09	0.473858007
	100 bars	0.000972959	3.71646E-09	1.36295E-09	4.44819E-09	0.465288554
ADAUSDT	50 bars	0.001736012	8.62496E-09	1.27048E-08	4.90547E-09	0.475664534
	100 bars	0.001543807	7.83887E-09	4.24928E-09	4.90575E-09	0.468764089
BTCUSDT	50 bars	30.59170154	2.22919E-05	42.21893443	4.85441E-10	0.477695275
	100 bars	27.11981393	2.00616E-05	1.100938512	4.85445E-10	0.469784298
ETHUSDT	50 bars	2.275569375	2.63985E-06	0.005914603	9.99954E-10	0.483894069
	100 bars	2.006449812	2.37222E-06	0.00496139	9.99933E-10	0.476930492

Table 10: Mean values of market microstructure variables after Nov 2021

Type	Window	Roll Measure	Roll Impact	Kyle's Lambda	Amihud	VPIN
SOLUSDT	50 bars	0.059758002	1.17637E-06	9.44109E-05	3.00358E-08	0.466968126
	100 bars	0.044615633	1.07671E-06	6.04667E-05	3.07251E-08	0.458688473
XRPUSDT	50 bars	0.000356229	6.30279E-09	3.17064E-09	1.13931E-08	0.460511516
	100 bars	0.00029473	5.84671E-09	1.02199E-08	1.18039E-08	0.451379484
ADAUSDT	50 bars	0.000605024	3.49444E-08	2.89657E-09	4.55577E-08	0.458147685
	100 bars	0.000554756	3.25889E-08	3.02232E-10	4.42423E-08	0.45164024
BTCUSDT	50 bars	14.28409348	1.52287E-05	0.391117916	3.94517E-10	0.46421365
	100 bars	12.83488862	1.39385E-05	1.347228563	3.94512E-10	0.453986963
ETHUSDT	50 bars	1.267135794	2.94883E-06	0.000162075	1.31292E-09	0.462581119
	100 bars	1.132305714	2.70517E-06	0.004874742	1.31298E-09	0.452947326

Figures

Figure 1: Average MDA own feature importance for change of realized volatility.

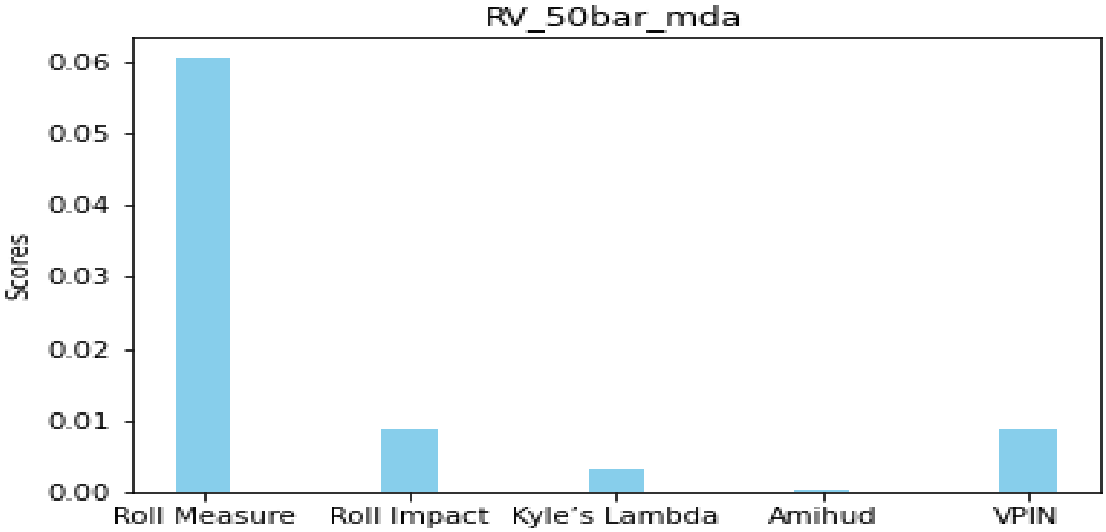

MDA feature importance using only own features for change of realized volatility. The prediction window is 50 bars. Results are aggregated across all five cryptocurrencies.

Figure 2: Average MDA feature importance for change of Kurtosis.

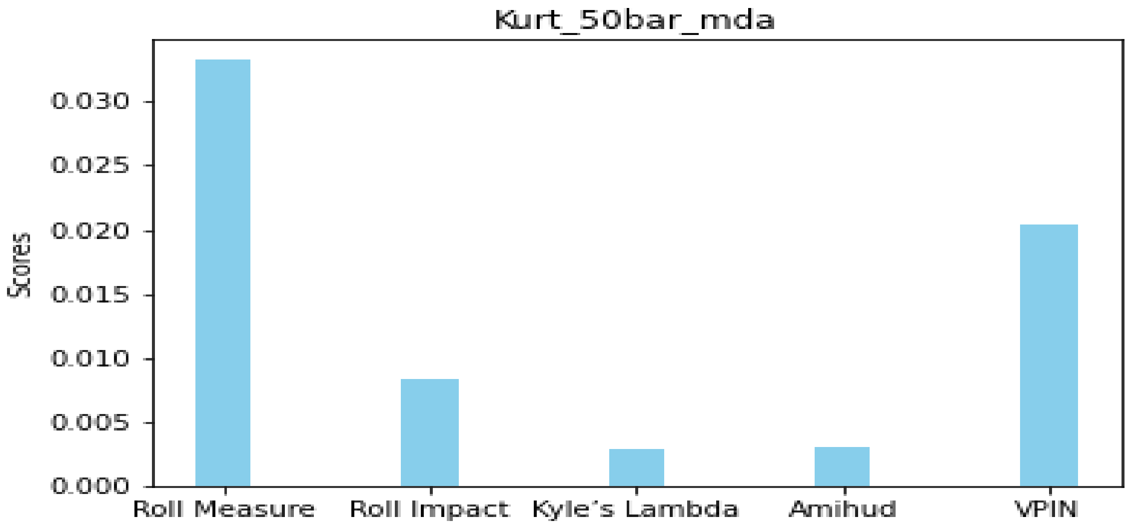

MDA feature importance using only own features for the change of Kurtosis. The prediction window is 50 bars. Results are aggregated across all five cryptocurrencies

Figure 3: MDA feature importance for change of JB Statistics

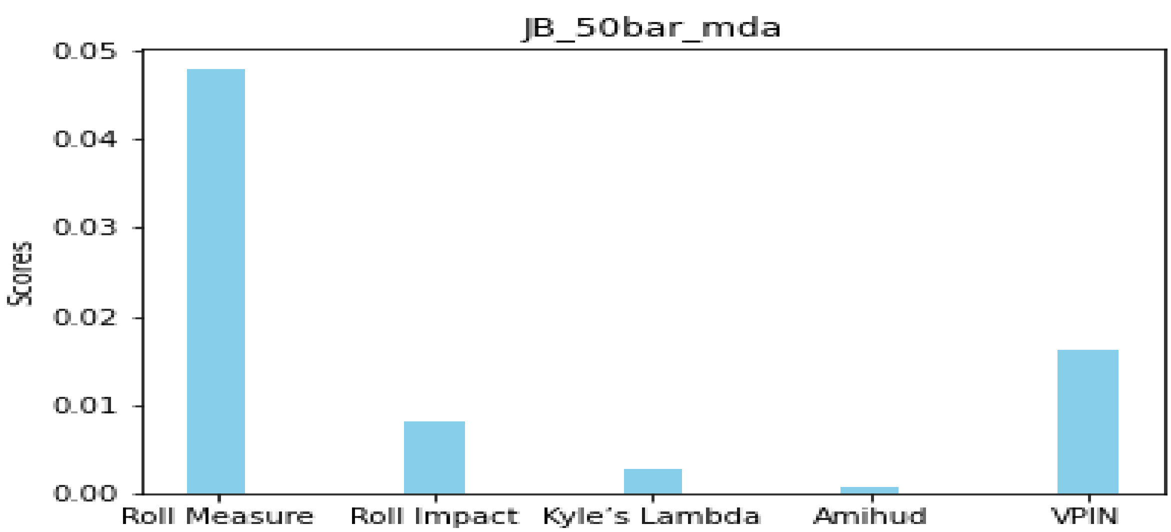

MDA feature importance using only own features for the change in the JB statistic. The prediction window is 50 bars. Results are aggregated across all five cryptocurrencies.

Figure 4: MDA feature importance for change of return auto-correlation.

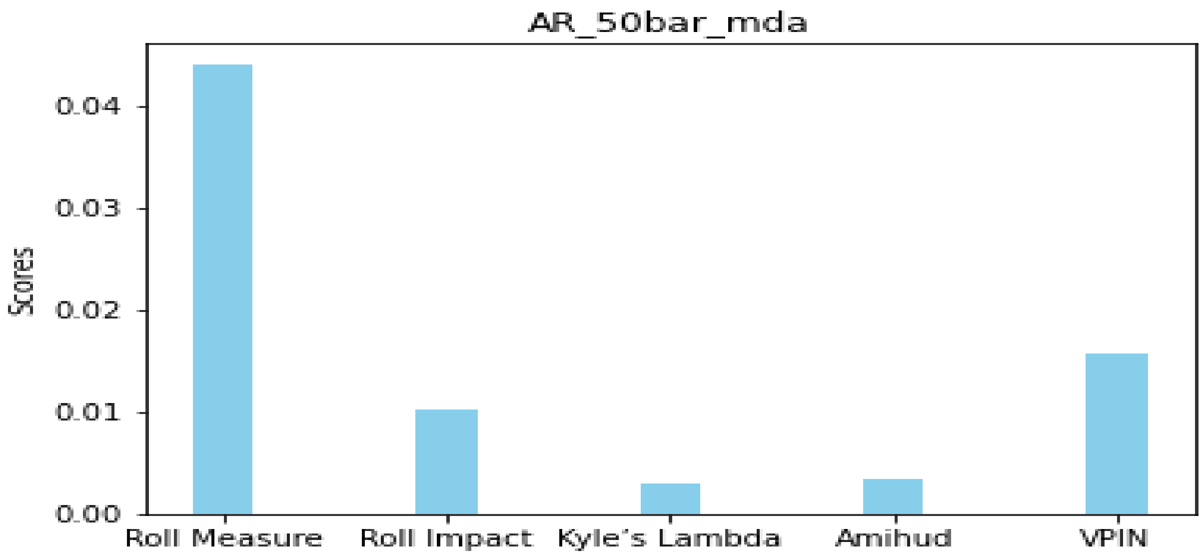

MDA feature importance using only own features for the change in return autocorrelation. The prediction window is 50 bars. Results are aggregated across all five cryptocurrencies.

Figure 5: Aggregated MDA scores for ADA using all 25 features

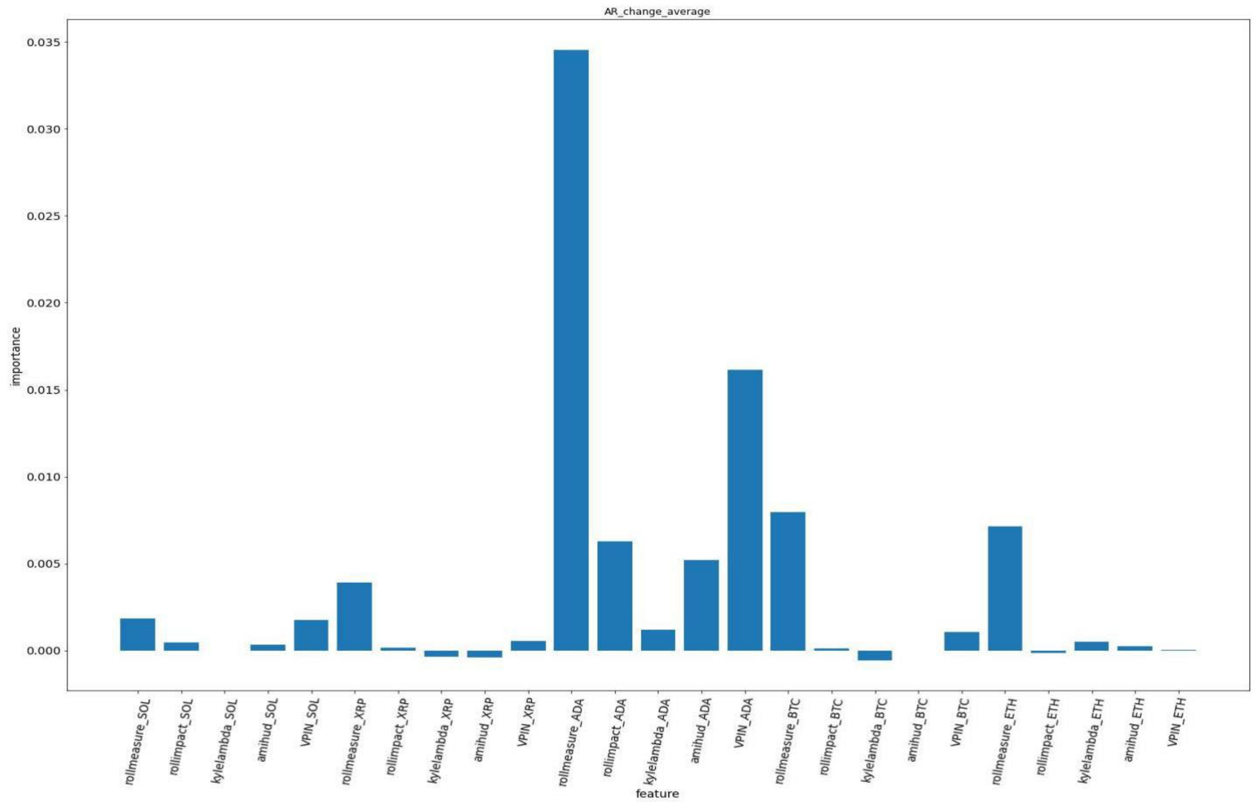

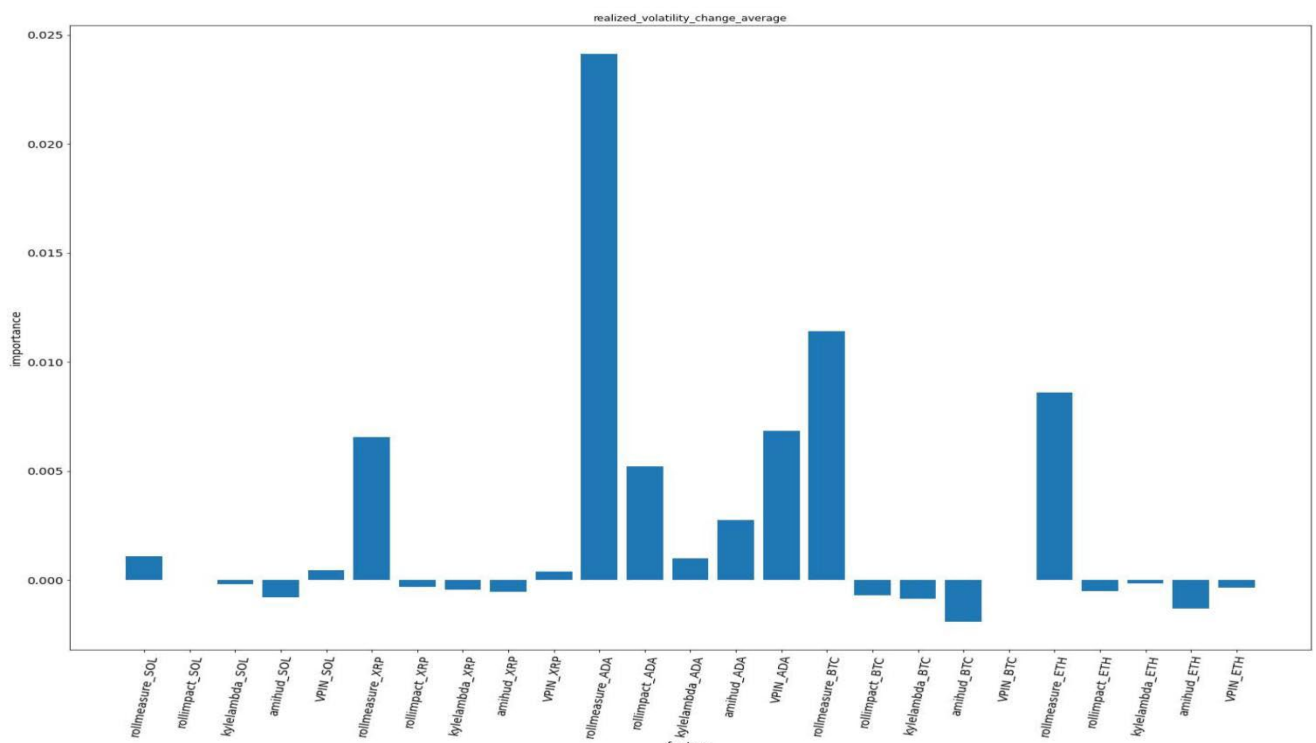

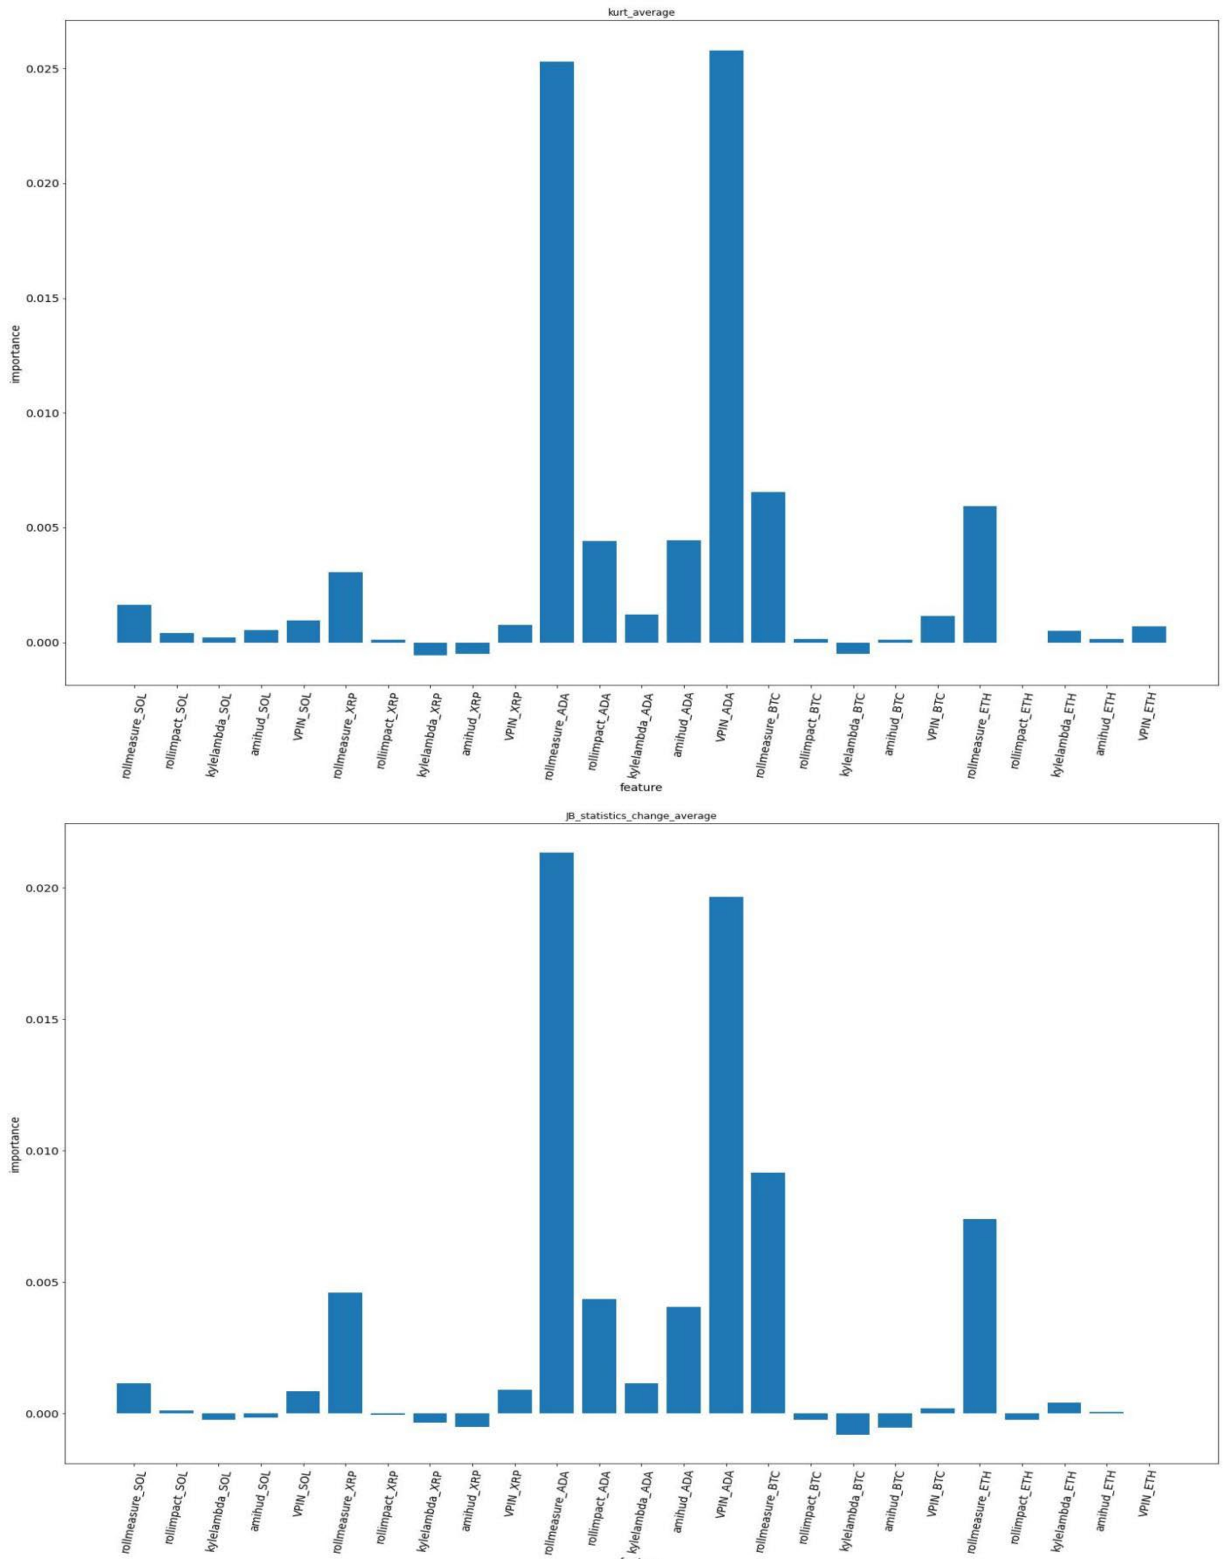

Each panel of this figure provides MDA scores (aggregated over 50 and 100 bars) for prediction of market statistics for ADA using all 25 features (five features for each of the five crypto currencies).

Figure 6: Aggregated MDA scores for BTC using all 25 features

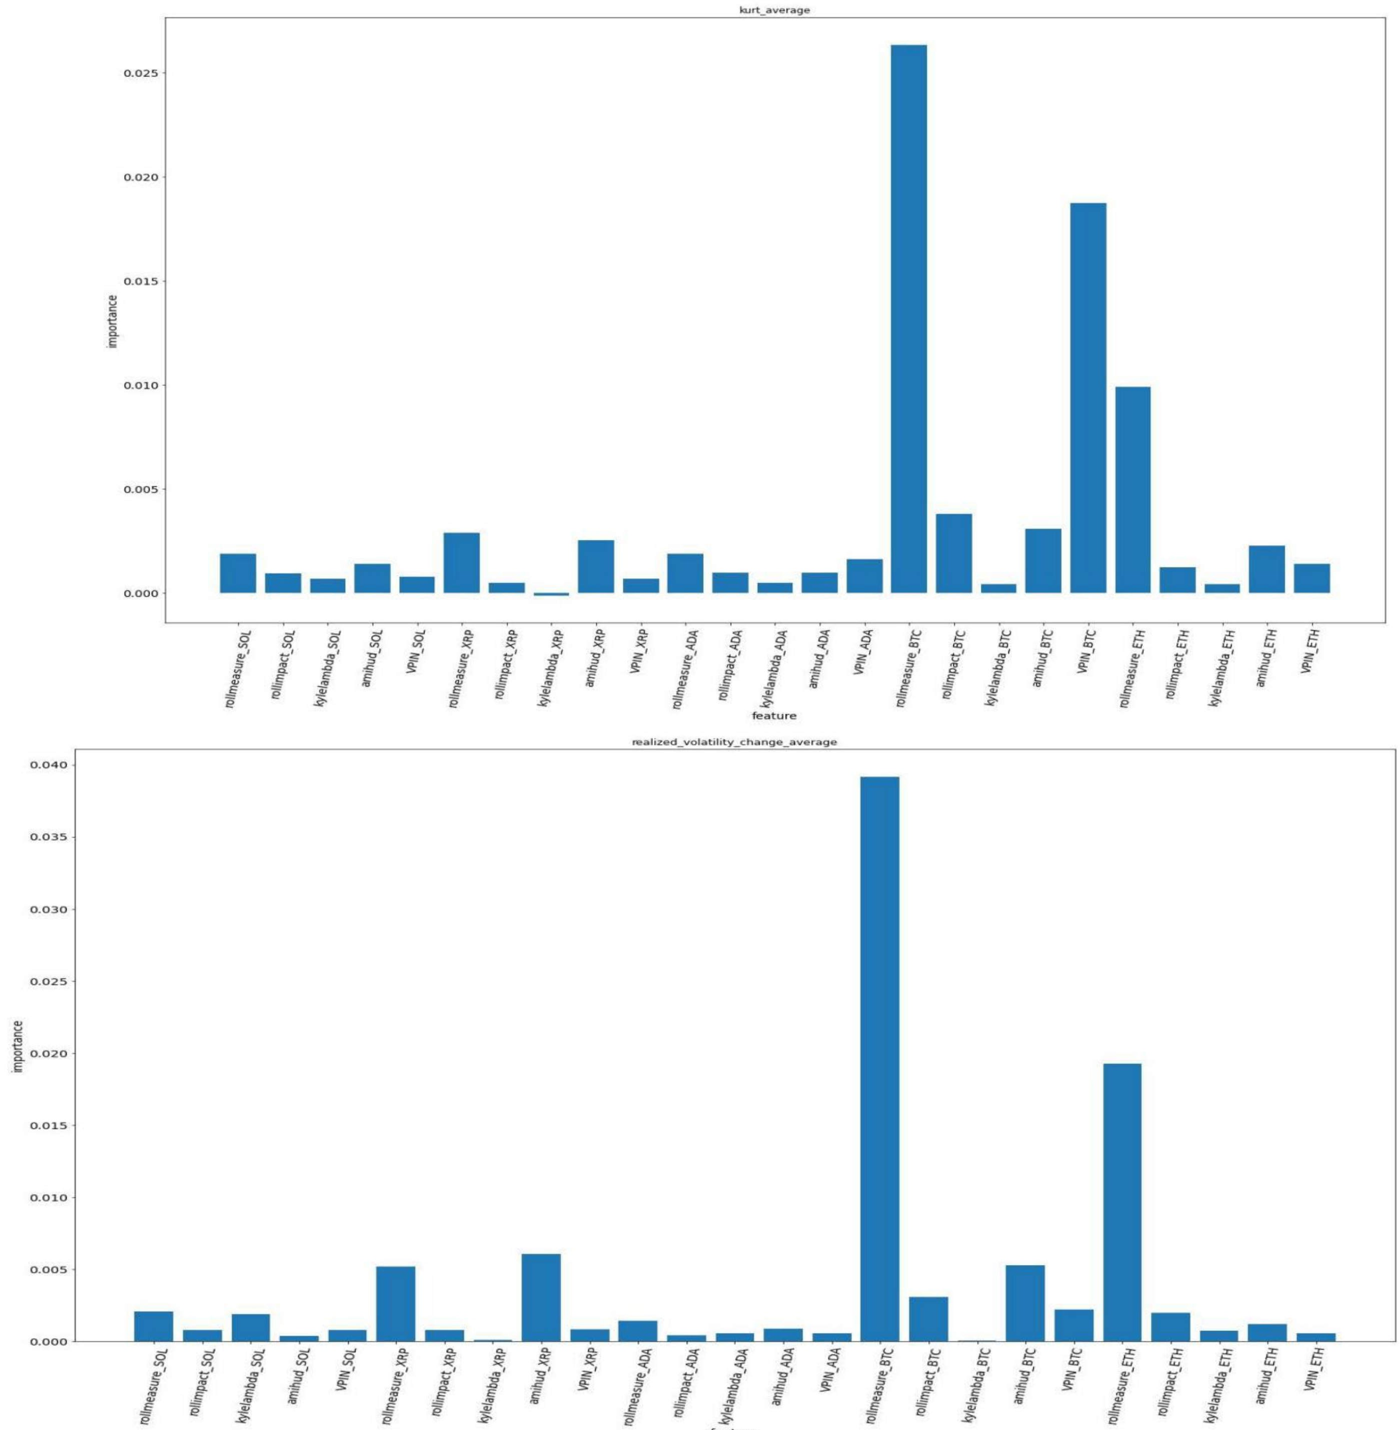

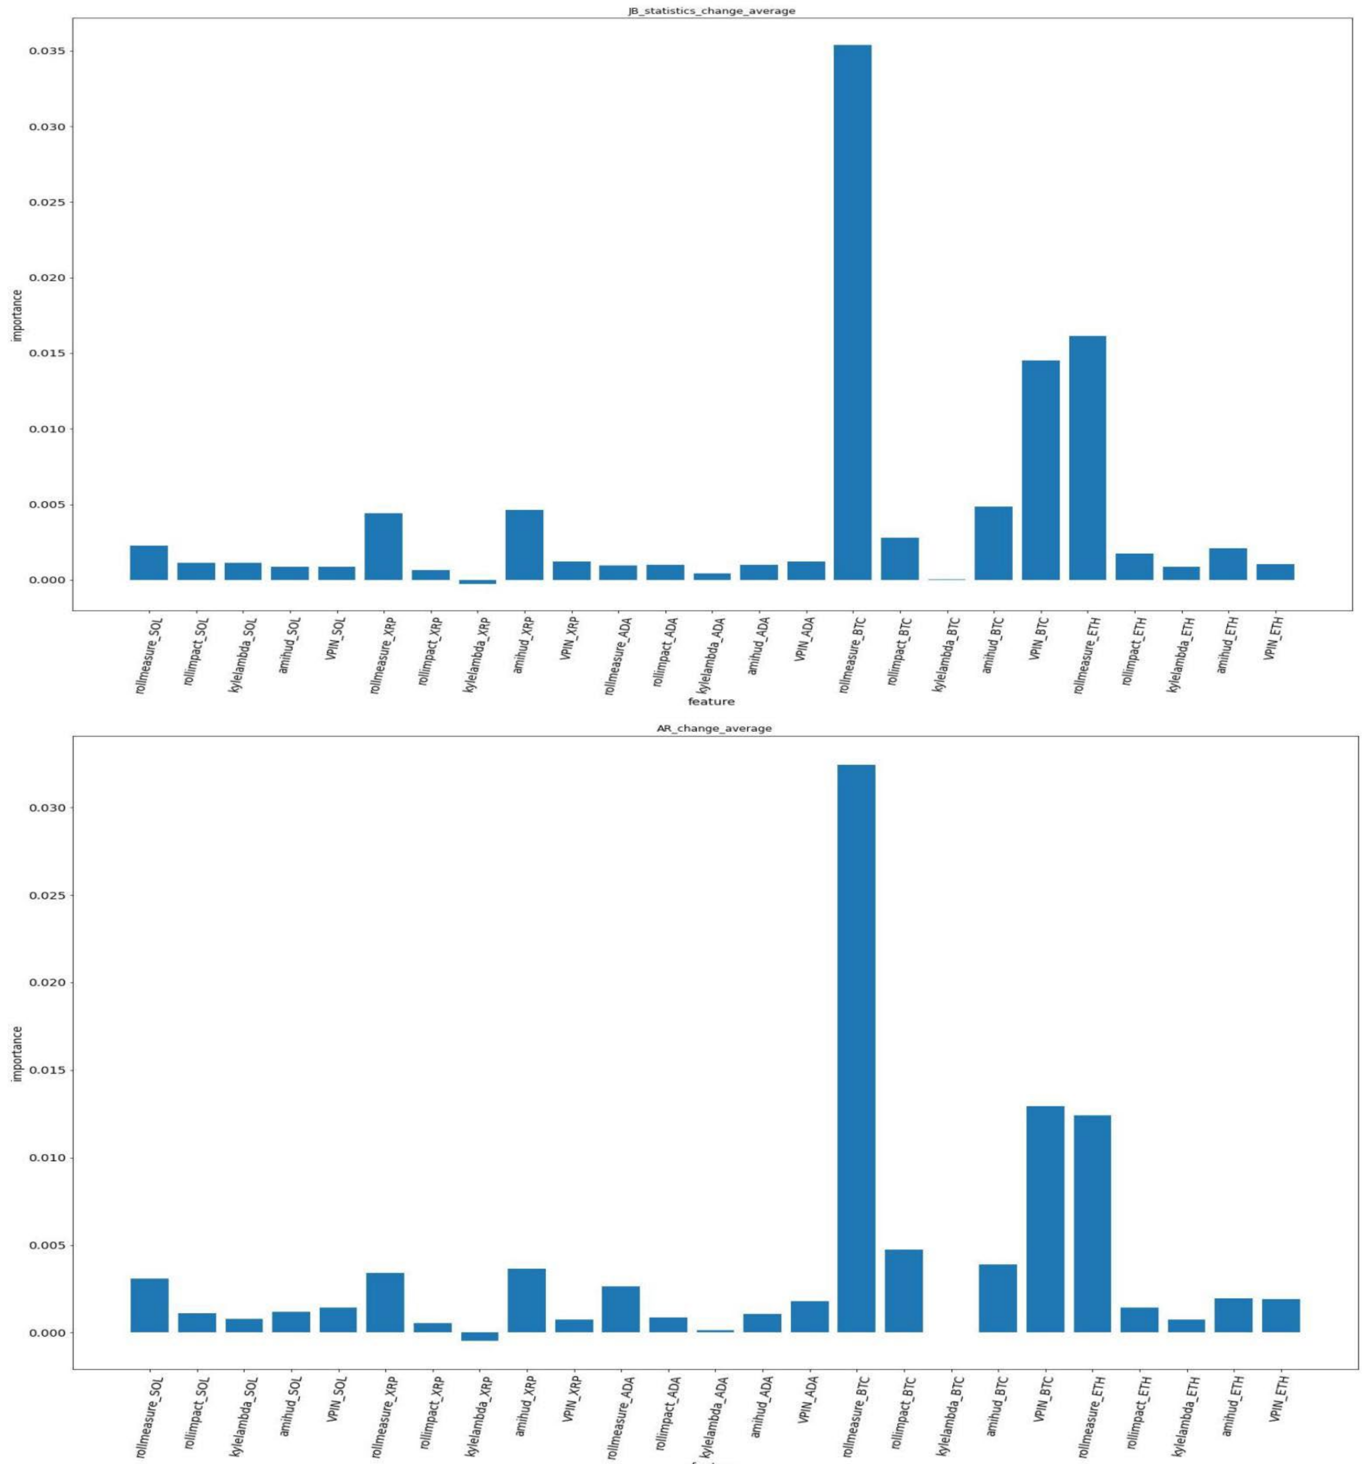

Each panel of this figure provides MDA scores (aggregated over 50 and 100 bars) for prediction of market statistics for BTC using all 25 features (five features for each of the five crypto currencies).

Figure 7: Aggregated MDA scores for ETH using all 25 features

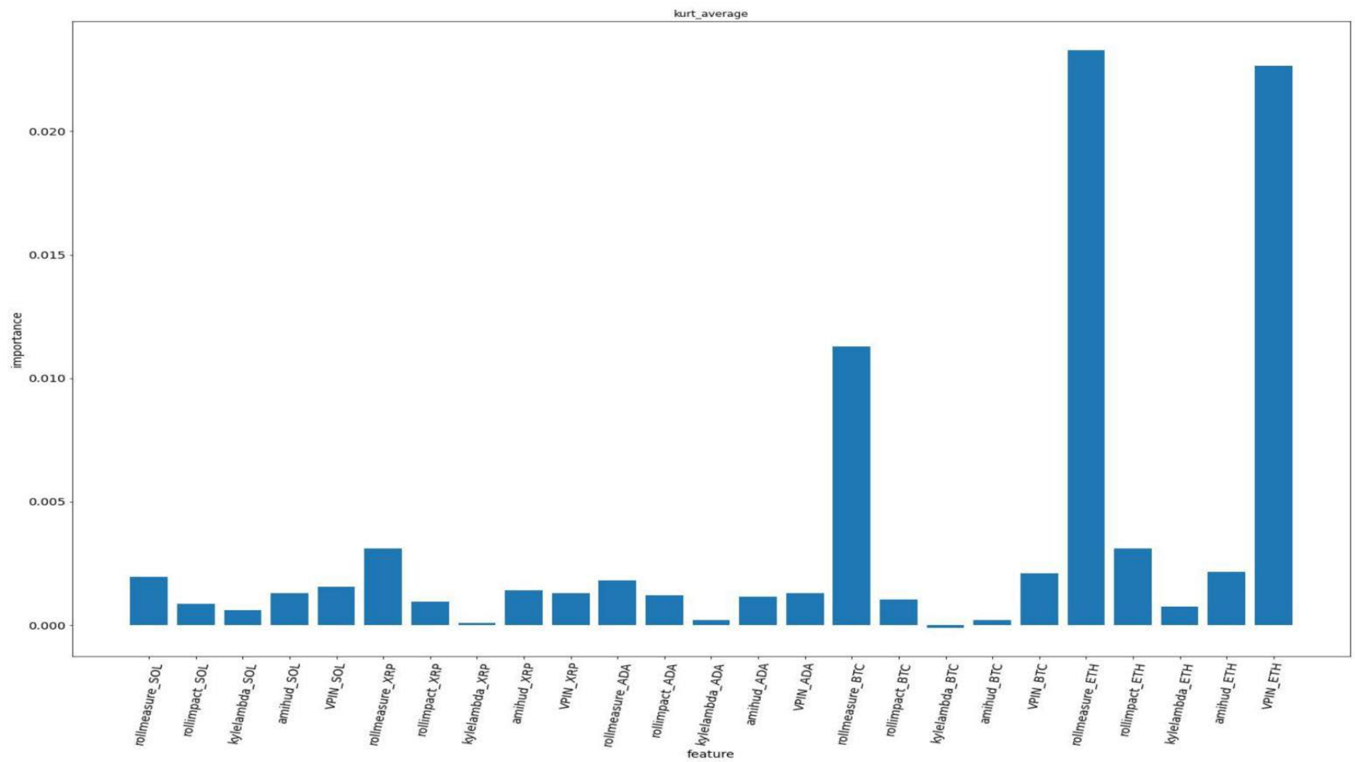

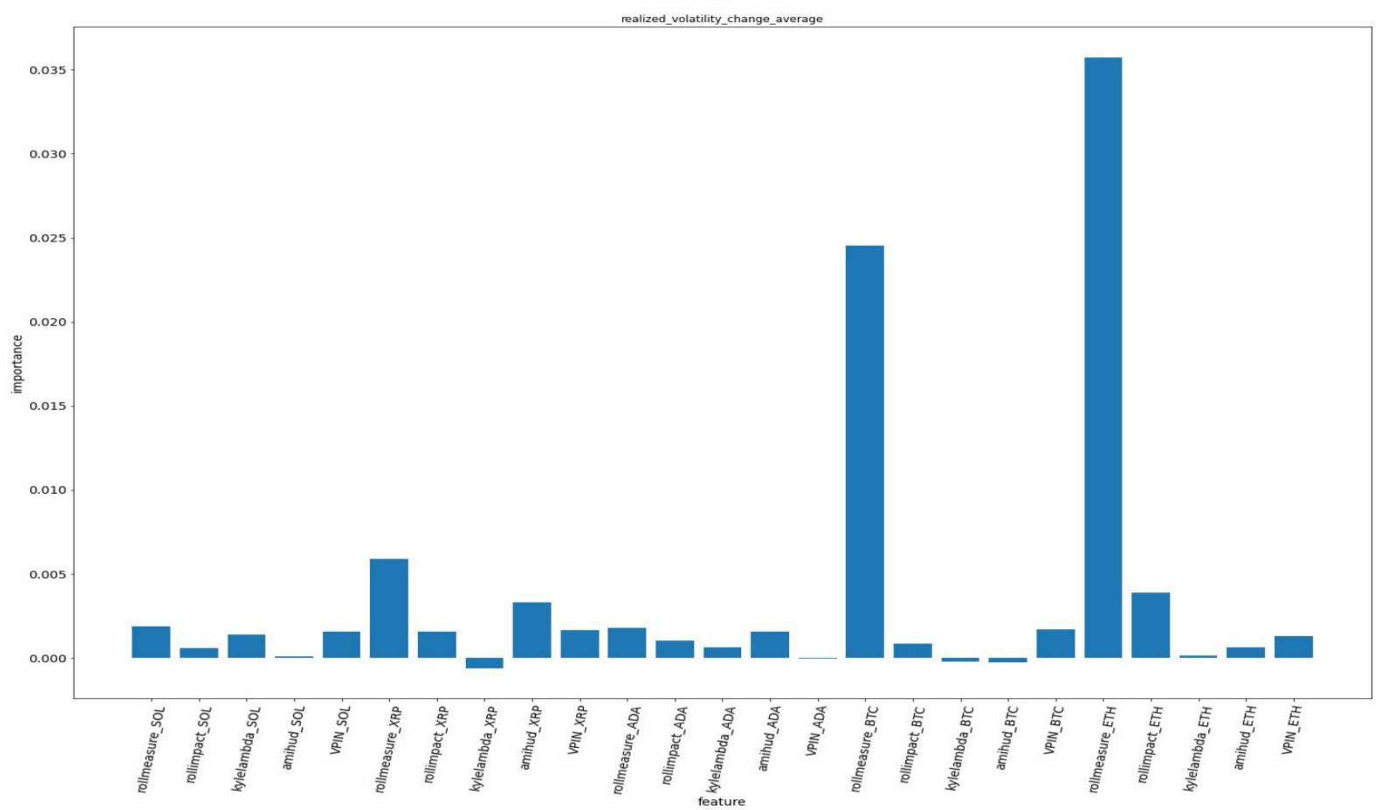

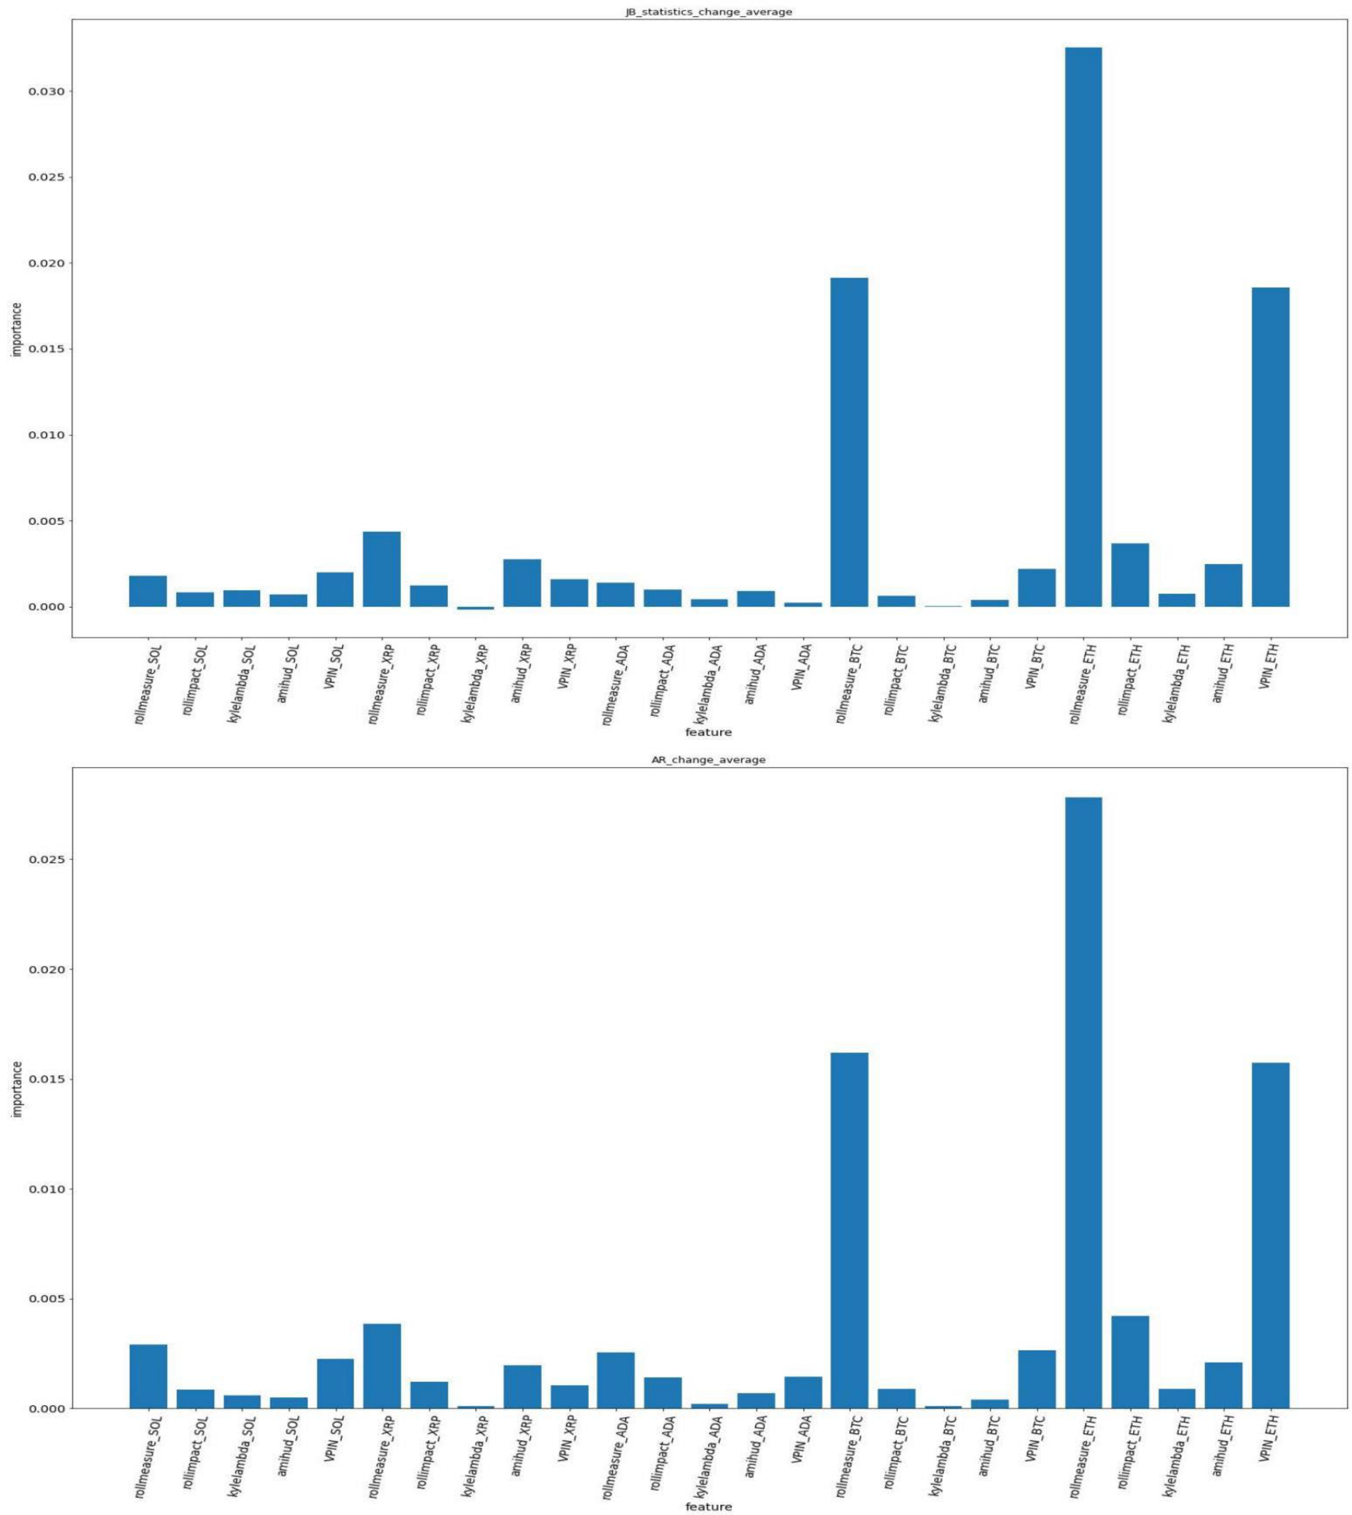

Each panel of this figure provides MDA scores (aggregated over 50 and 100 bars) for prediction of market statistics for ETH using all 25 features (five features for each of the five crypto currencies).

Figure 8: Aggregated MDA scores for SOL using all 25 features

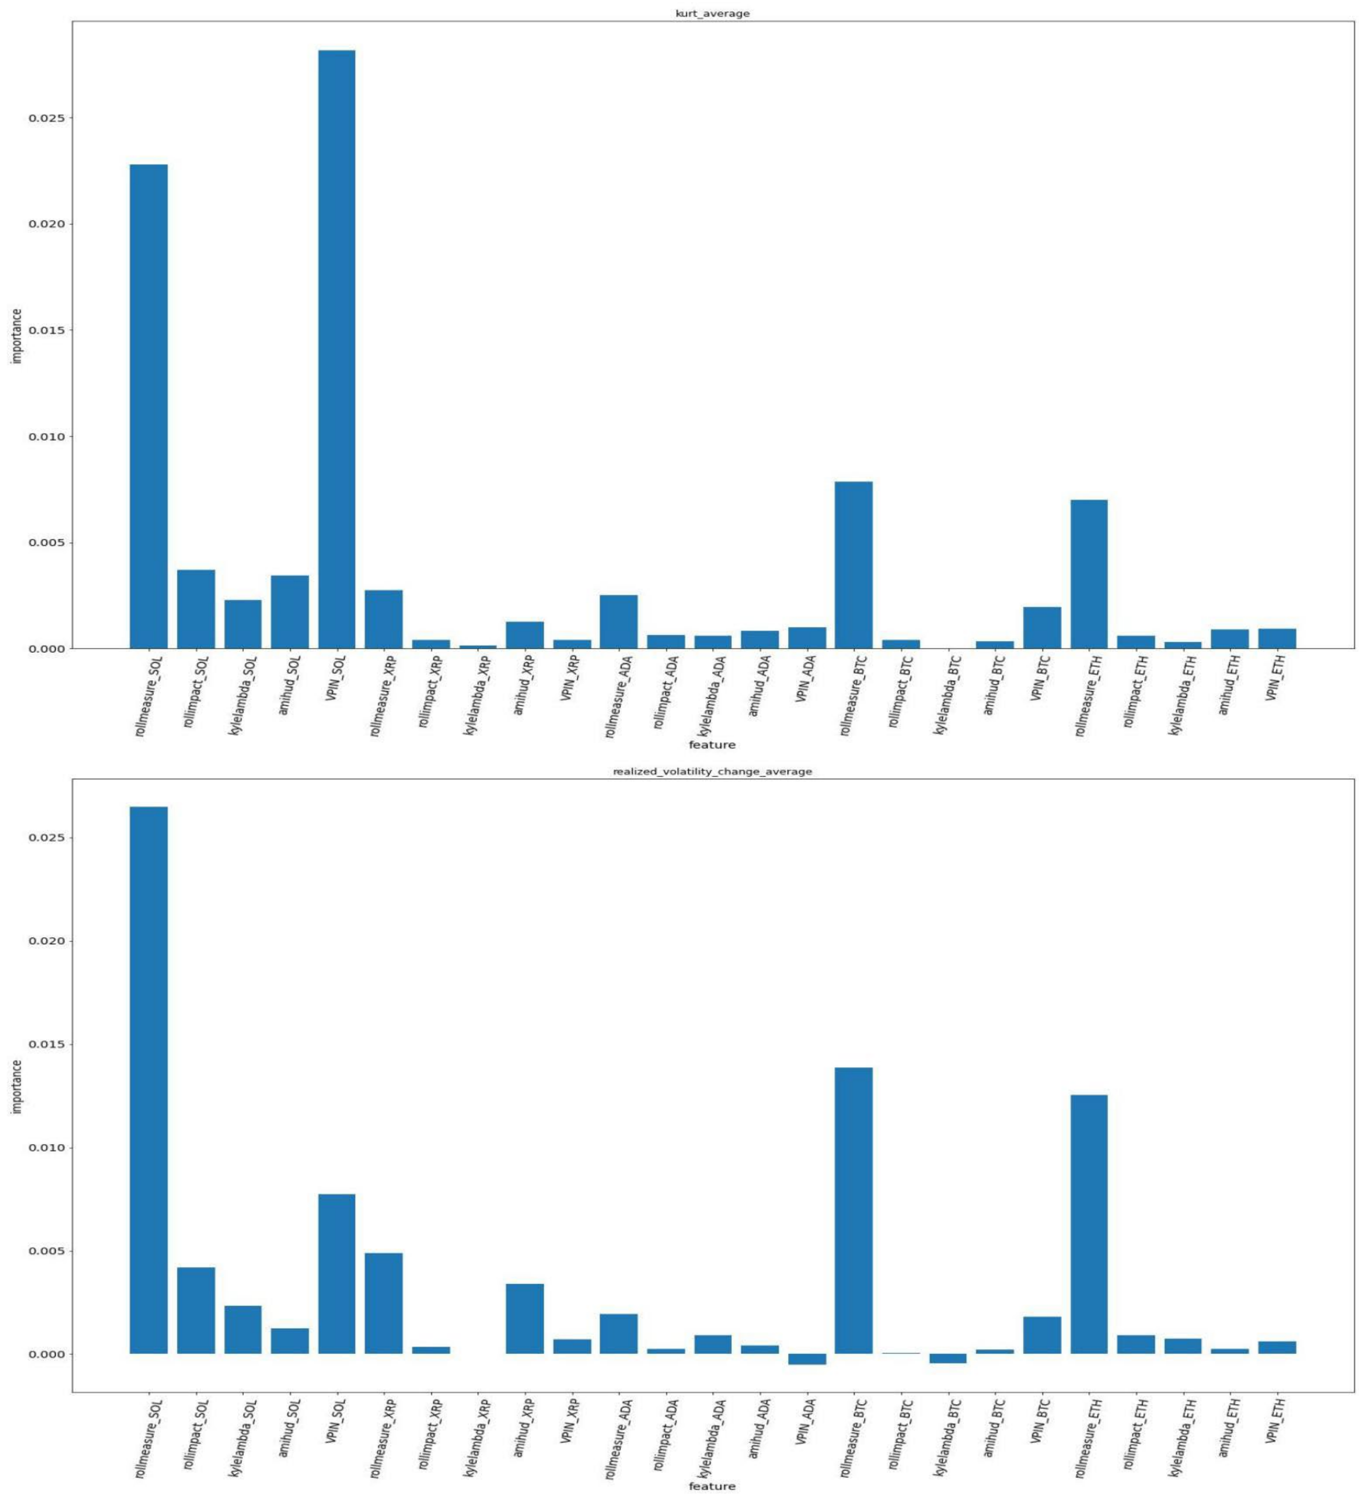

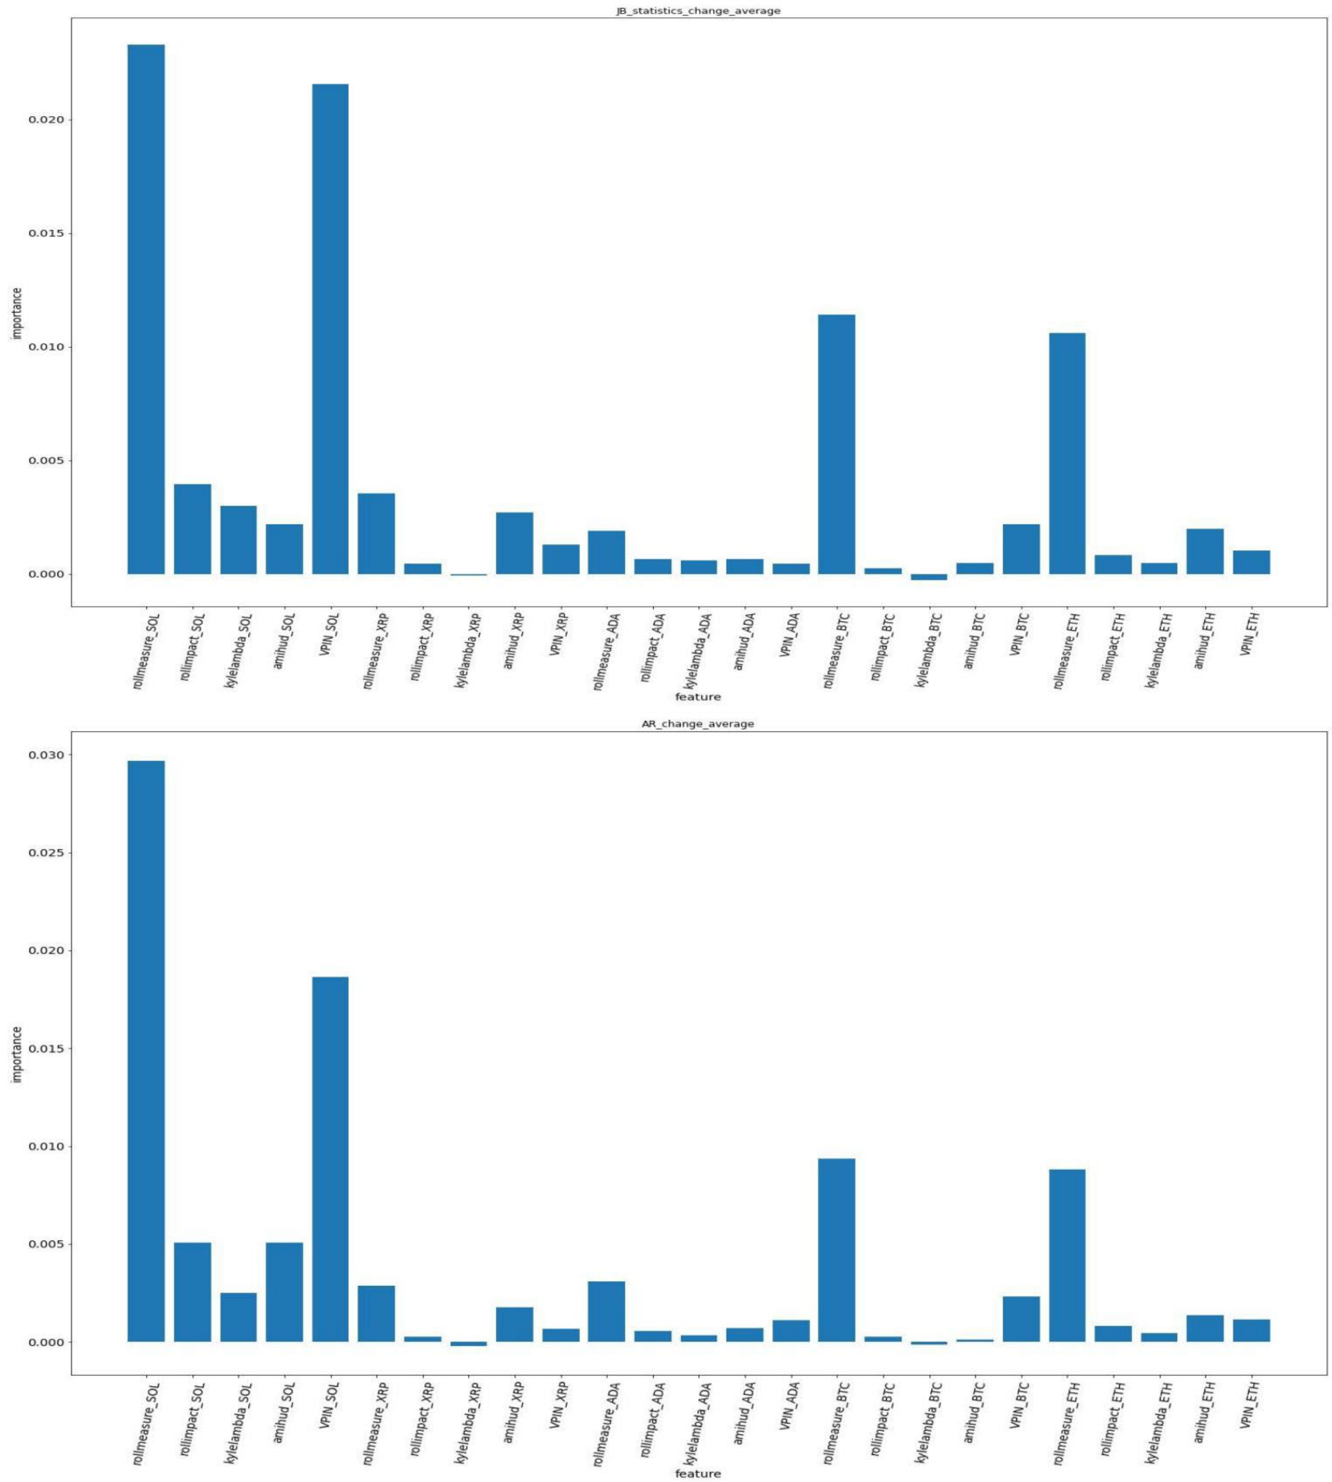

Each panel of this figure provides MDA scores (aggregated over 50 and 100 bars) for prediction of market statistics for SOL using all 25 features (five features for each of the five crypto currencies).

Figure 9: Aggregated MDA scores for XRP using all 25 features

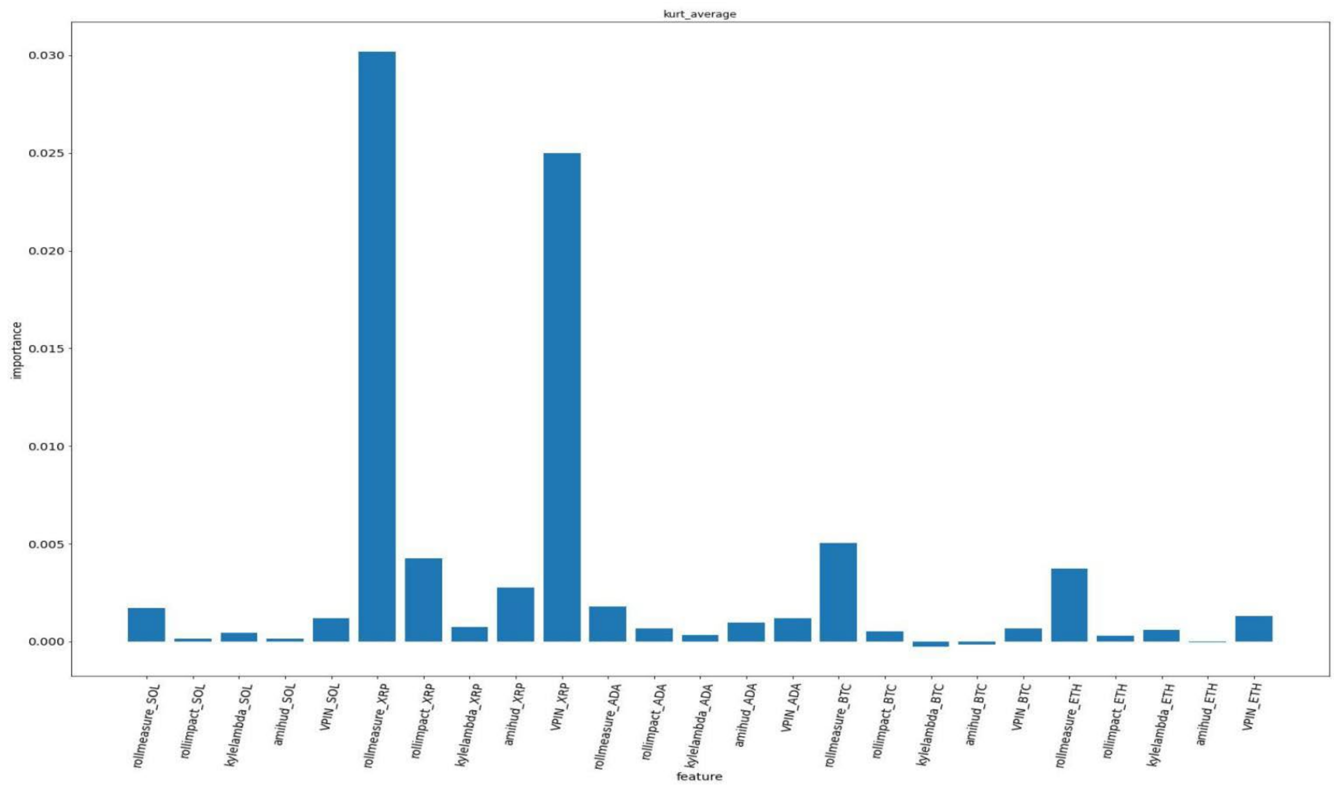

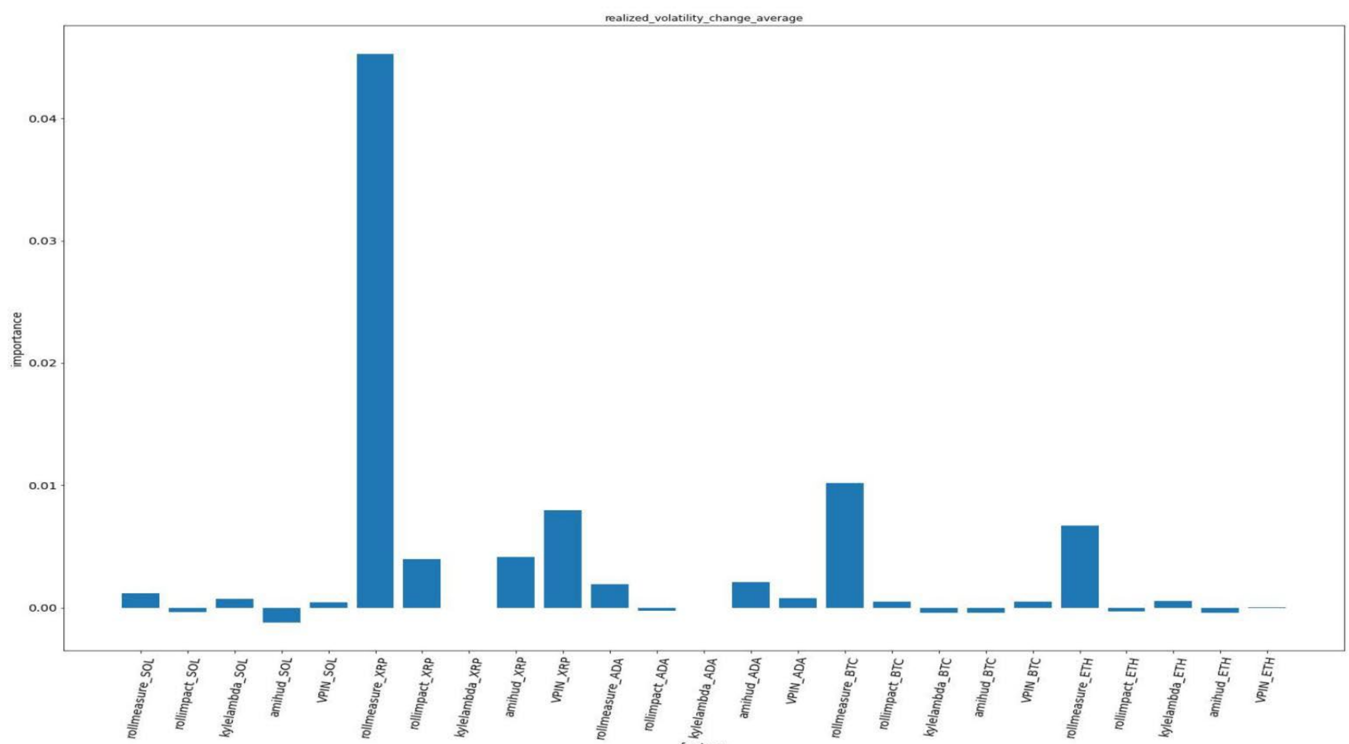

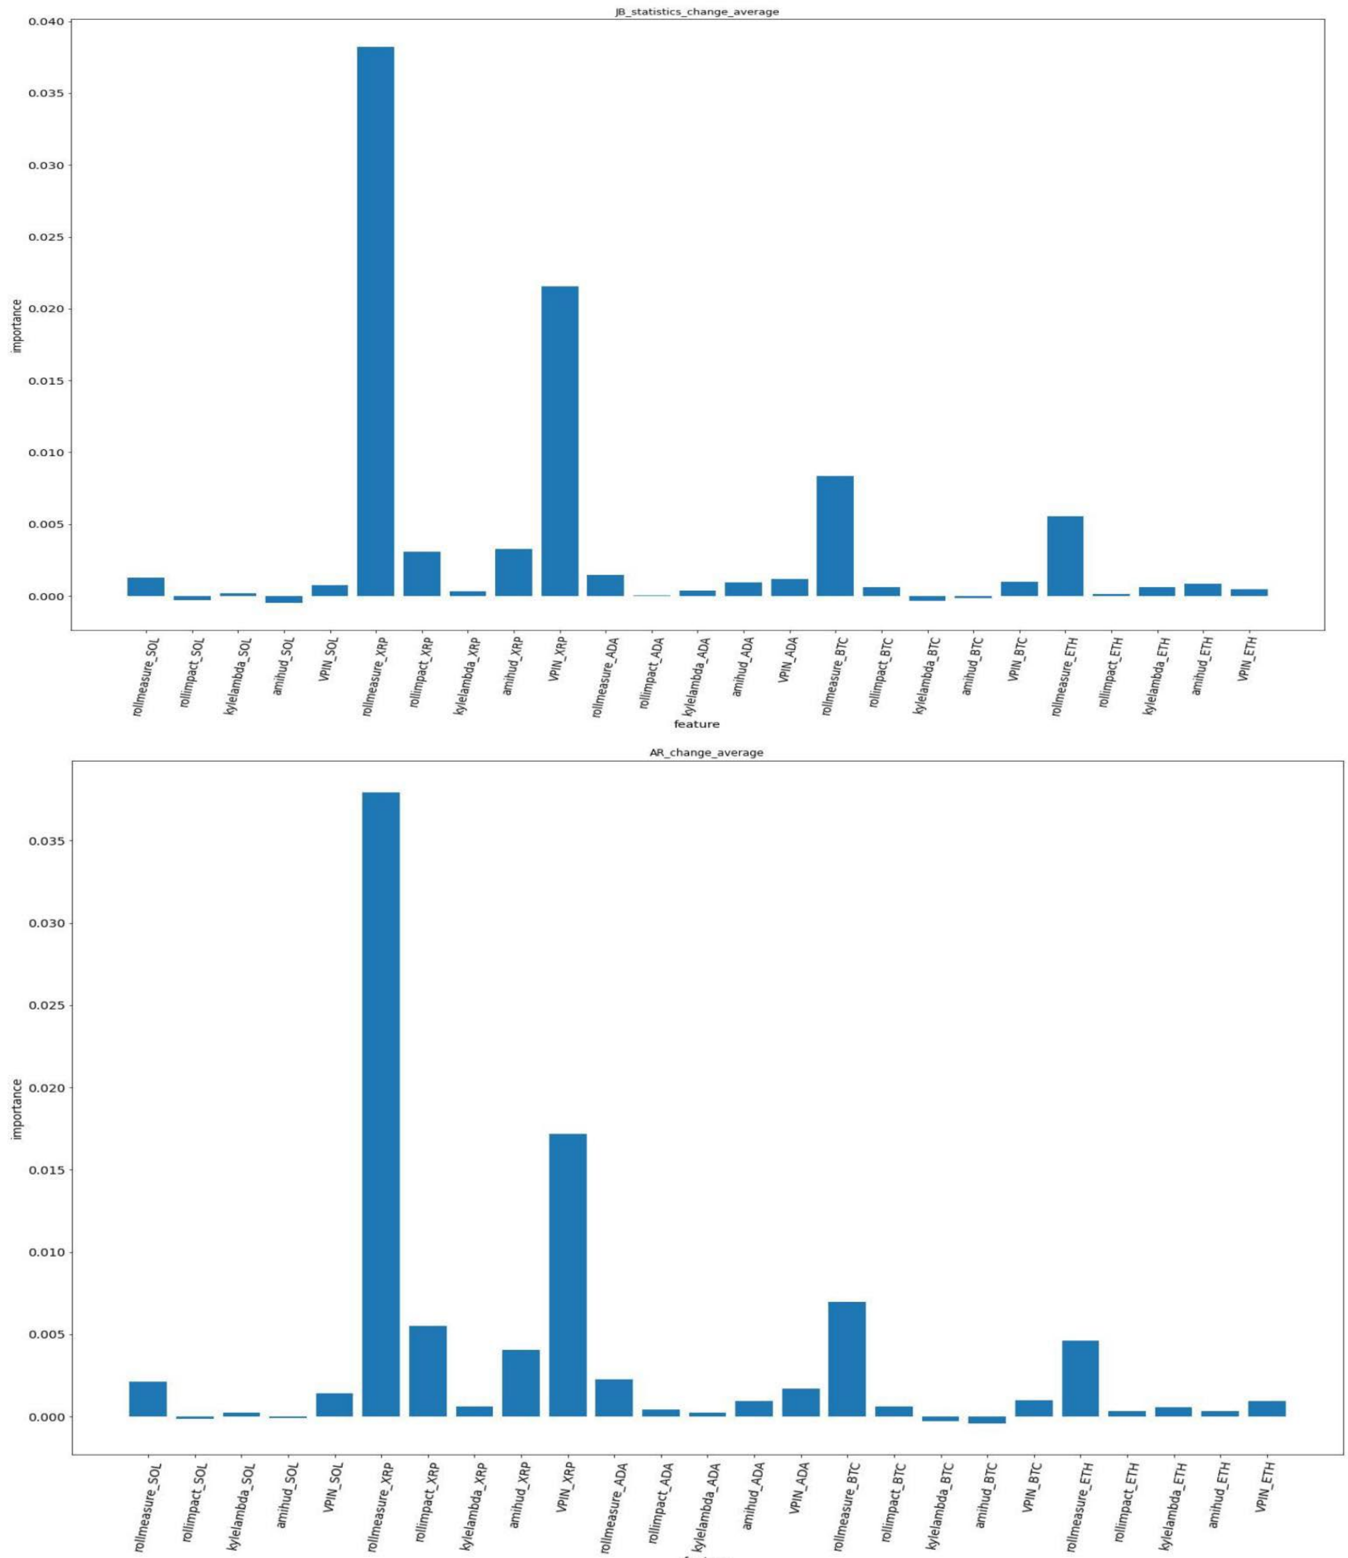

Each panel of this figure provides MDA scores (aggregated over 50 and 100 bars) for prediction of market statistics for XRP using all 25 features (five features for each of the five crypto currencies).

Figure 10: Aggregated Accuracy and MDS Results from a Logistic Regression

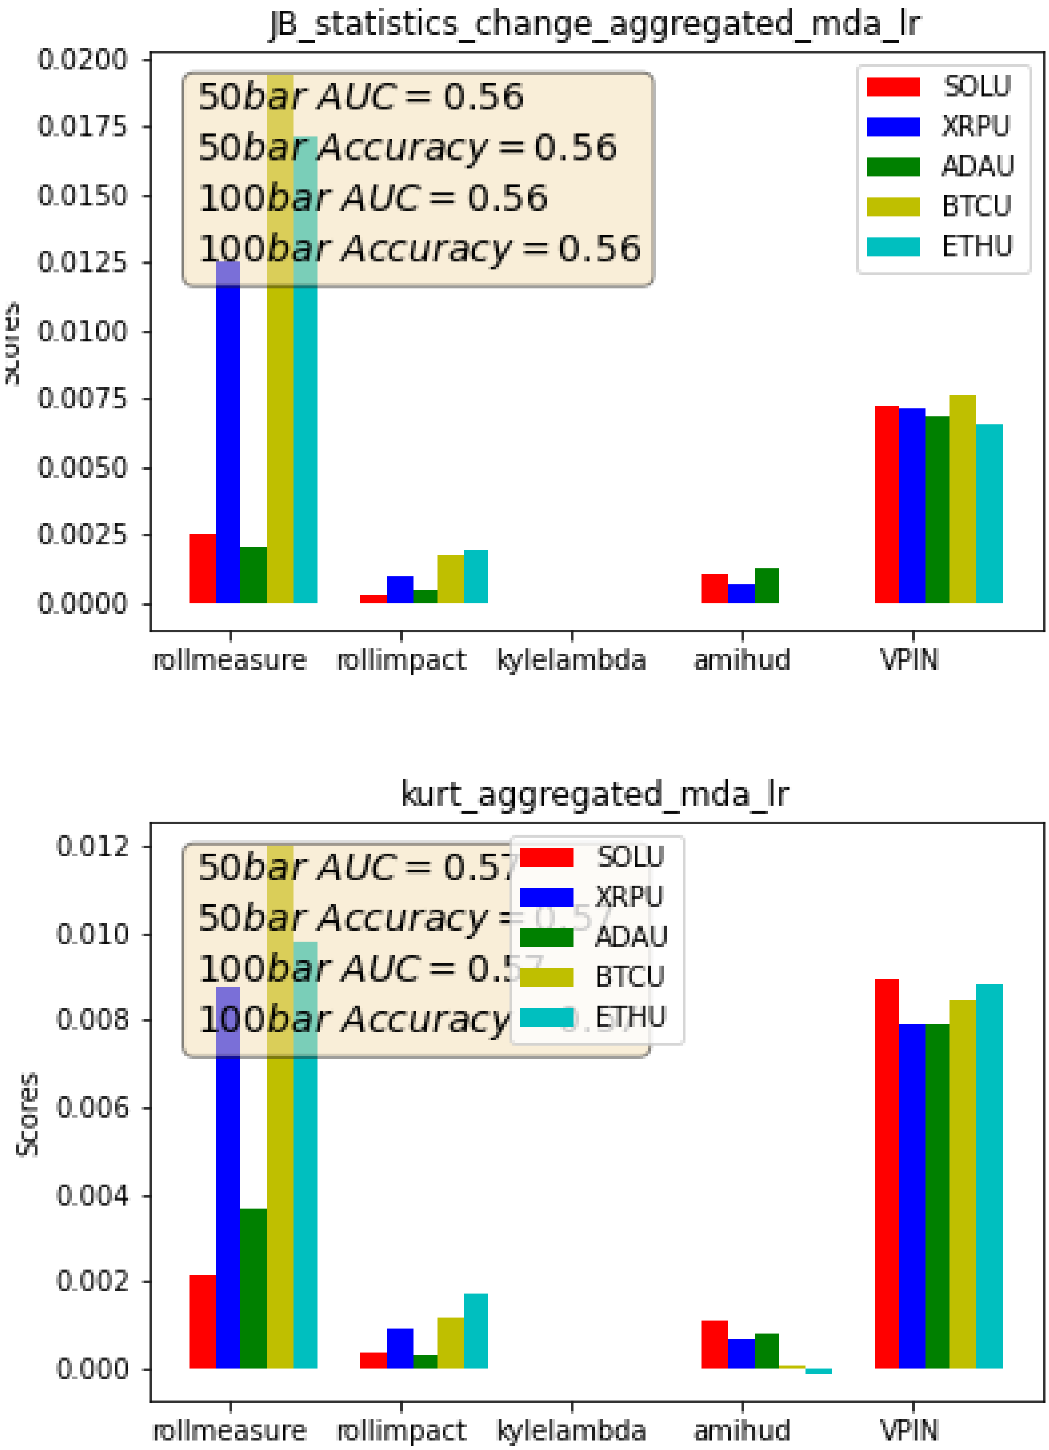

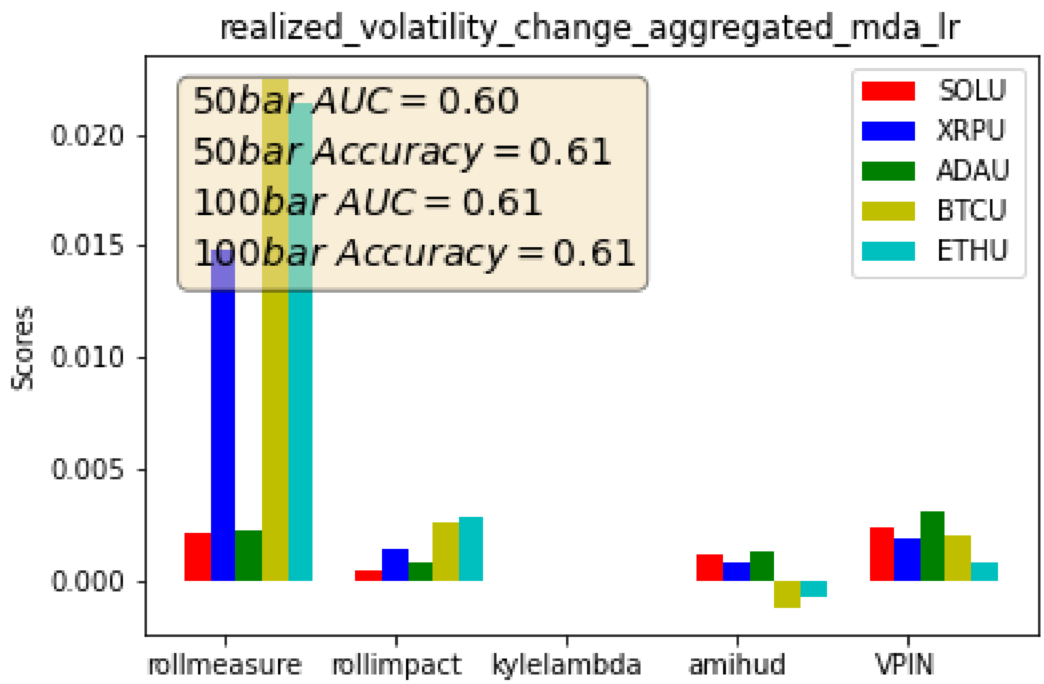
realized_volatility_change_aggregated_mda_lr
$50bar\ AUC = 0.60$
$50bar\ Accuracy = 0.61$
$100bar\ AUC = 0.61$
$100bar\ Accuracy = 0.61$
SOLU
XRPU
ADAU
BTCU
ETHU
Scores
0.020
0.015
0.010
0.005
0.000
rollmeasure
rollimpact
kylelambda
amihud
VPIN

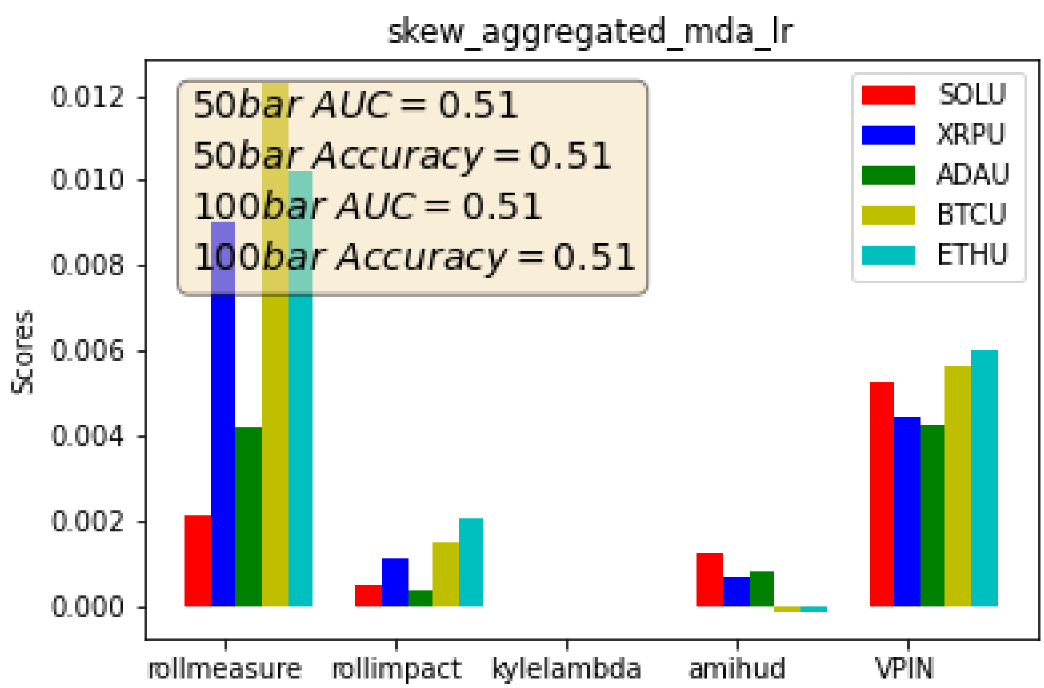
skew_aggregated_mda_lr
$50bar\ AUC = 0.51$
$50bar\ Accuracy = 0.51$
$100bar\ AUC = 0.51$
$100bar\ Accuracy = 0.51$
SOLU
XRPU
ADAU
BTCU
ETHU
Scores
0.012
0.010
0.008
0.006
0.004
0.002
0.000
rollmeasure
rollimpact
kylelambda
amihud
VPIN