PREDICTABILITY OF VOLATILITY OF CRYPTOCURRENCIES

Young Li
Bloomberg Enterprise Quants
{xli528, entquant}@bloomberg.net

Tuesday 28^{th} December, 2021

ABSTRACT

Prices of cryptocurrencies have shown large fluctuations. However, they have been progressively adopted as an asset class by financial institutions. A question naturally arises: is the crypto class exceptional in terms of its volatility dynamics and predictability? In this study, we have selected instruments from a wide range of categories for comparison with cryptocurrencies. We show that from both the perspectives of persistence and mean reversion, cryptocurrencies is similar to instruments in other asset classes. We then examine volatility predictability using the heterogeneous autoregressive (HAR) model and Random Forests (RFs). In-sample and out-of-sample model performances are evaluated using the correlation coefficient between observed and predicted. In both cases, the correlation decreases with increasing prediction horizon. The out-of-sample correlation drops more rapidly than the in-sample case, and the RF model in general outperforms the HAR model. From the correlation-prediction horizon curves, we see that the crypto asset class behaves in line with instruments in other asset classes. We conclude that the volatility of cryptocurrencies is on average as predictable as for other instruments.

1 Are Cryptocurrencies Exceptional?

The prices of cryptocurrencies have shown large fluctuations. In 2016, the price of Bitcoin rose by 125% and in 2017 the price rose by more than 2,000%. Following the 2017 peak, Bitcoin's price receded once more. In 2021, Bitcoin continued to set new all-time highs, more than tripling the peak price Bitcoin achieved during the 2017 bull run. Following the illegalization of all cryptocurrency transactions or facilitation announced by China's government and central bank, Bitcoin's price dropped more than half through August 2021 to around $29,700 as miners scrambled to relocate (Reiff 2022).

Why are cryptocurrencies so volatile? Apart from the government regulation factor, the traditional supply-and-demand mechanism is pointed out as a dominant factor (Reiff 2022): by design, Bitcoin is limited to 21 million coins–the closer the circulating supply gets to this limit, the higher prices are likely to climb. But it is difficult to predict what will happen to prices when the limit is reached. On another note, cryptocurrencies are still new and remain in the price discovery phase when their "terminal values" are still undefined (Sigalos 2021).

Cryptocurrencies has been progressively adopted as an asset class by financial institutions. According to a 2021 annual global crypto hedge fund report by PwC (PwC 2021), the total assets under management (AuM) of crypto hedge funds globally increased to nearly US$3.8 billion in 2020 from US$2 billion the previous year, and the percentage of crypto hedge funds with AuM over US$20 million increased in 2020 from 35% to 46%. According to the 2021 Q3 quarterly report published by Crypto Fund Research (CryptoFundResearch 2021), as of the end of report period, there were more than 850 crypto funds across the globe with primary offices in more than 80 countries, and the total crypto fund assets under management climbed to US$59.6 billion from US$8.3 billion in 2018 Q3.

There are well-known models for describing the dynamics of volatility and for forecasting, such as ARCH (Engle 1982), GARCH (Bollerslev 1986), HAR (Corsi 2009), etc. A question naturally arises: Is the crypto asset class unique or exceptional in terms of its volatility dynamics and predictability?

In order to assess uniqueness or characteristics of crypto volatility, we have selected instruments from a wide range of categories for comparison. Table 1 lists considered instruments grouped by categories (*Others* are used not for comparison but as model input features in a later section).

The evolution of volatility over time is complex, but from the vast literature on this subject, two properties are widely recognized: persistence and mean reversion. We will discuss both in the following subsections.

1.1 Persistence

An important property of volatility recognized widely in literature is persistence (e.g., Mandelbrot 1963, Ding et al. 1993, Cont et al. 1997, and Guillaume et al. 1997). A quantitative manifestation of it is that, while returns themselves are uncorrelated, absolute returns or their squares (proxies for volatility) display a significant and slowly decaying autocorrelation (e.g., Ding et al. 1993, Cont 2007, etc.).

Figure 1 shows the partial autocorrelation functions (PACF) of absolute return timeseries for Bitcoin/USD cross (XBTUSD) along with five instruments representing currency, equity index, fixed income, commodity, and gold. The PACF gives the partial correlation of a stationary time series with its own lagged values, effectively removing the effects of shorter lags. The use of this function was introduced by Box et al. (2008) for identifying the order of an autoregressive process. We use PACF here as the manifestation of self persistence. XBTUSD has a close to 0.4 lag-1 value, comparable to the high values for SPX, GT10 and CL1. It takes about 10 days for the XBTUSD PACF to fall below a low reference level of 0.05, again comparable to other instruments.

1.2 Mean Reversion

Mean reversion is another generally observed property of volatility (Papanicolaou and Sircar 1999). It's formulated in several stochastic volatility models, two widely referenced of which are examined with our data:

- Feller or Cox-Ingersoll-Ross process (Cox et al. 1985, Heston 1993):

$$\mathrm{d}\nu_t = \alpha(\mu - \nu_t)\mathrm{d}t + \sigma\sqrt{\nu_t}\mathrm{d}W_t \tag{1.1}$$

- Ornstein-Uhlenbeck process (Nelson 1990, Papanicolaou and Sircar 1999):

$$d\nu_t = \alpha(\mu - \nu_t)dt + \sigma dW_t \tag{1.2}$$

where ν_t is the instantaneous variance, W_t is a Wiener process, and α, μ and σ are parameters. The parameter α corresponds to the speed of adjustment to the mean μ, and σ corresponds to volatility.

Table 1: Instruments in major asset classes/categories used for comparison of volatility characteristics.

Asset Class	Instrument	Description
Crypto	XBTUSD Curncy	Bitcoin/USD Cross Rate
	XBNUSD Curncy	Bitcoin Cash/USD Cross Rate
	XETUSD Curncy	Ethereum/USD Cross Rate
	XDHUSD Curncy	DASH/USD Cross Rate
	EOSUSD Curncy	EOS/USD Cross Rate
	XTHUSD Curncy	Ethereum Classic/USD Cross Rate
	XLCUSD Curncy	Litecoin/USD Cross Rate
	XMRUSD Curncy	Monero/USD Cross Rate
	XRPUSD Curncy	XRP/USD Cross Rate
	XZCUSD Curncy	Zcash/USD Cross Rate
Equity Index	SPX Index	S&P 500 Index
	UKX Index	FTSE 100 Index
	NKY Index	Nikkei 225 Index
	MXWO Index	MSCI World Index
	MXEF Index	MSCI Emerging Markets Index
	MXWD Index	MSCI ACWI Index
Fixed Income	GT10 Govt	Generic United States 10 Year Government Note
	GTFN10Y Govt	Generic FANNIE MAE 10 Year Government Bond
	GTFH10Y Govt	Generic FREDDIE MAC 10 Year Government Bond
	GTJPY10Y Govt	Generic Japan 10 Year Government Bond
	GTDEM10Y Govt	Generic Germany 10 Year Government Bond
Commodity	CL1 Comdty	Generic 1st Crude Oil, WTI
	NG1 Comdty	Generic 1st Natural Gas
	HG1 Comdty	Generic 1st Copper
	W 1 Comdty	Generic 1st Wheat
	CT1 Comdty	Generic 1st CT Future
Currency	EURUSD Curncy	EUR/USD
	GBPUSD Curncy	GBP/USD
	AUDUSD Curncy	AUD/USD
	USDJPY Curncy	USD/JPY
	DXY Curncy	US Dollar Index Spot Rate
	XAU Curncy	Gold United States Dollar Spot
Crypto-Related Equity	RIOT US Equity	Riot Blockchain Inc (U.S.)
	MARA US Equity	Marathon Digital Holdings Inc (U.S.)
	GBTC US Equity	Grayscale Bitcoin Trust BTC (U.S.)
Others	M2 Index	Federal Reserve United States Money Supply M2 SA
	VIX Index	Chicago Board Options Exchange Volatility Index

Figure 2 plots increments of squared weekly returns against their starting values for XBTUSD along with five other instruments. We use the squared return here as a proxy for variance. All six instruments show a decreasing trend, from positive increments when starting values are low to negative increments when starting values are high. The trends are also approximately linear, consistent with Equation 1.1 and Equation 1.2 with $\mathrm{d}t$ equal to one week (similar trends exist for daily increments with slightly larger scattering). This inverse relationship makes volatility eventually revert to its long-run average level.

The volatility model parameters can be estimated using maximum-likelihood. For the CIR process, $\nu_{t+\Delta}$ (Δ being a small increment) follows a generalized chi-square distribution given ν_t (Cox et al. 1985); for the Ornstein-Uhlenbeck process, $\nu_{t+\Delta}$ follows a Normal distribution given ν_t (Papanicolaou and Sircar 1999). Table 2 shows model parameters and negative log-likelihood estimated using weekly returns for the 10 cryptocurrencies and other 5 instruments. The negative log-likelihood (NLL) values for the CIR model are much smaller than those for the Ornstein-Uhlenbeck model, showing the former is a better model. Under the CIR model, the volatility σ parameters are much greater for the cryptocurrencies than for the other 5 instruments. But the noise levels relative to the long-run averages as measured by $\sigma/\sqrt{\mu}$ ratios are on average comparable.

Table 2: Volatility model parameters estimated using maximum-likelihood method. $r_w^2(t)/wd$ is used for ν_t in the models (Equation 1.1 and Equation 1.2), where $r_w(t)$ is weekly return at time t and wd number of days in a week (7 for cryptocurrencies and 5 for other instruments). NLL is the average negative log-likelihood at optimal model parameters. Data range: 2011-01 to 2021-12.

Instrument	Cox-Ingersoll-Ross process					Ornstein-Uhlenbeck process				
	α	μ	σ	$\sigma/\sqrt{\mu}$	NLL	α	μ	σ	σ/μ	NLL
XBTUSD Curncy	0.22	36.98	3.99	0.66	4.50	0.18	36.88	67.92	1.84	6.11
XBNUSD Curncy	0.47	49.66	6.86	0.97	4.90	0.40	49.75	101.01	2.03	6.15
XETUSD Curncy	1.43	31.51	9.49	1.69	4.45	1.02	31.51	83.55	2.65	5.49
XDHUSD Curncy	0.39	45.14	5.90	0.88	4.80	0.37	45.12	108.38	2.40	6.25
EOSUSD Curncy	1.34	35.91	9.81	1.64	4.58	0.91	35.88	106.34	2.96	5.79
XTHUSD Curncy	0.32	46.74	5.47	0.80	4.81	0.27	46.69	101.95	2.18	6.35
XLCUSD Curncy	1.48	31.18	9.61	1.72	4.44	1.25	31.18	83.59	2.68	5.39
XMRUSD Curncy	0.39	25.77	4.50	0.89	4.24	2.74	25.43	140.06	5.51	5.51
XRPUSD Curncy	0.51	41.47	6.52	1.01	4.72	0.38	41.37	95.18	2.30	6.11
XZCUSD Curncy	1.21	39.12	9.73	1.56	4.67	0.83	39.35	118.66	3.02	5.94
EURUSD Curncy	0.63	0.25	0.56	1.12	-0.40	0.56	0.25	0.43	1.75	0.52
SPX Index	0.50	0.89	0.94	1.00	0.86	0.22	0.89	1.80	2.01	2.36
GT10 Govt	0.39	6.92	2.32	0.88	2.88	0.18	6.92	10.86	1.57	4.22
CL1 Comdty	0.42	9.84	2.87	0.91	3.21	0.40	9.66	63.20	6.54	5.66
XAU Curncy	0.81	0.94	1.23	1.27	0.94	0.56	0.94	2.62	2.79	2.33

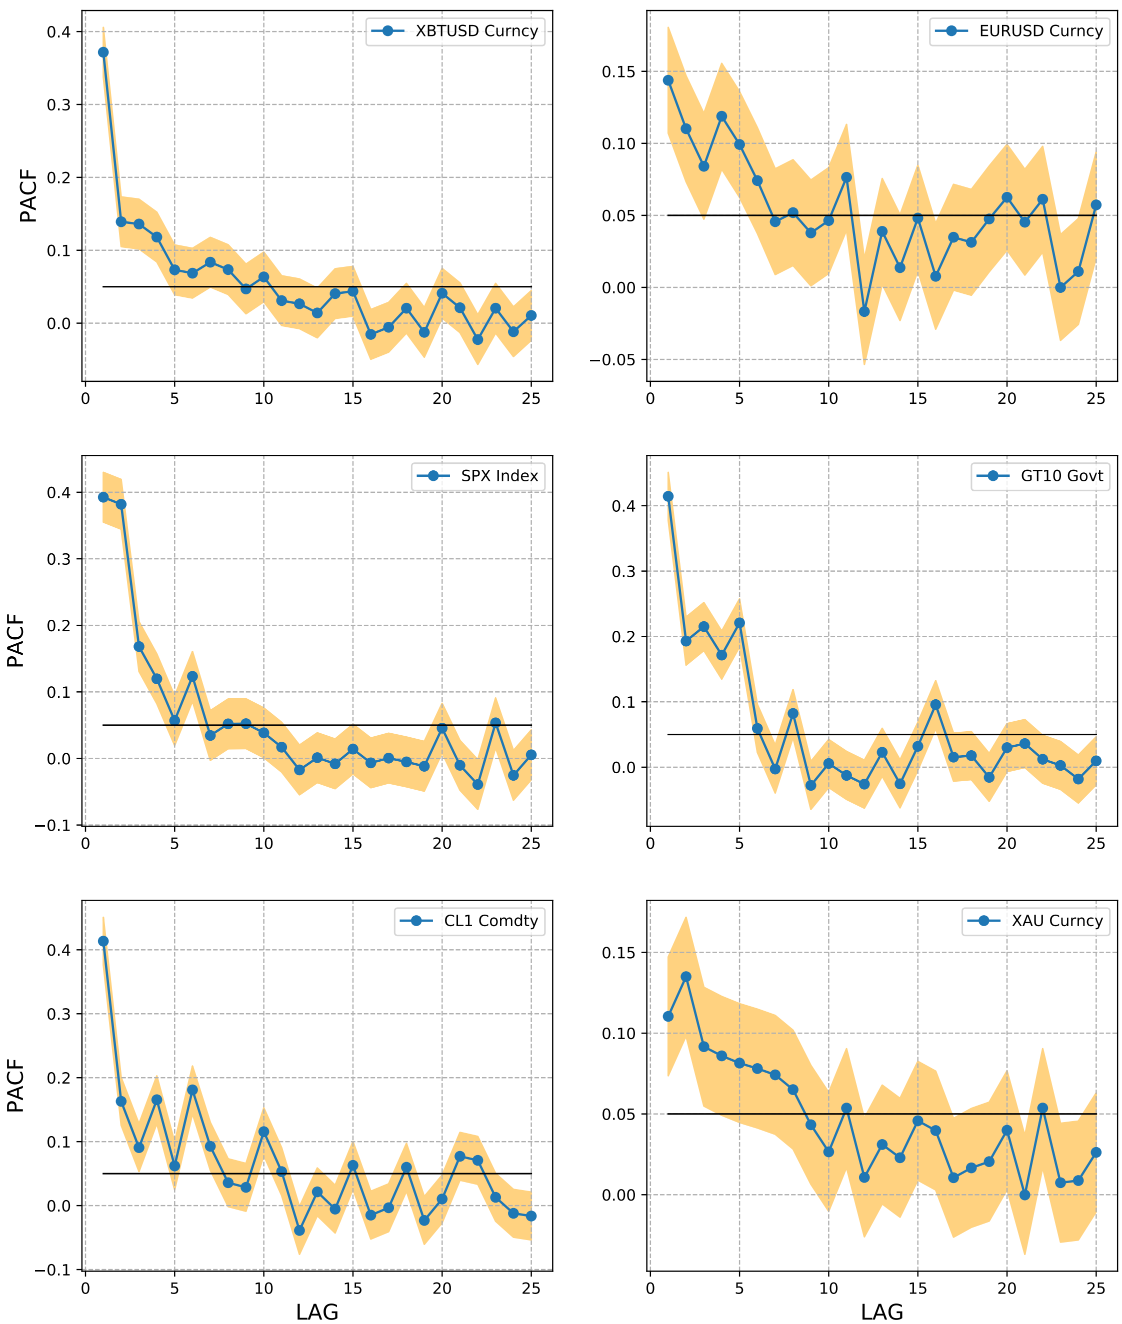

Figure 1: Partial autocorrelation functions of absolute return timeseries for six instruments. Shaded area indicates 95% confidence region. The black horizontal line in each subplot is for PACF = 0.05. Data range: 2011-01 to 2021-12.

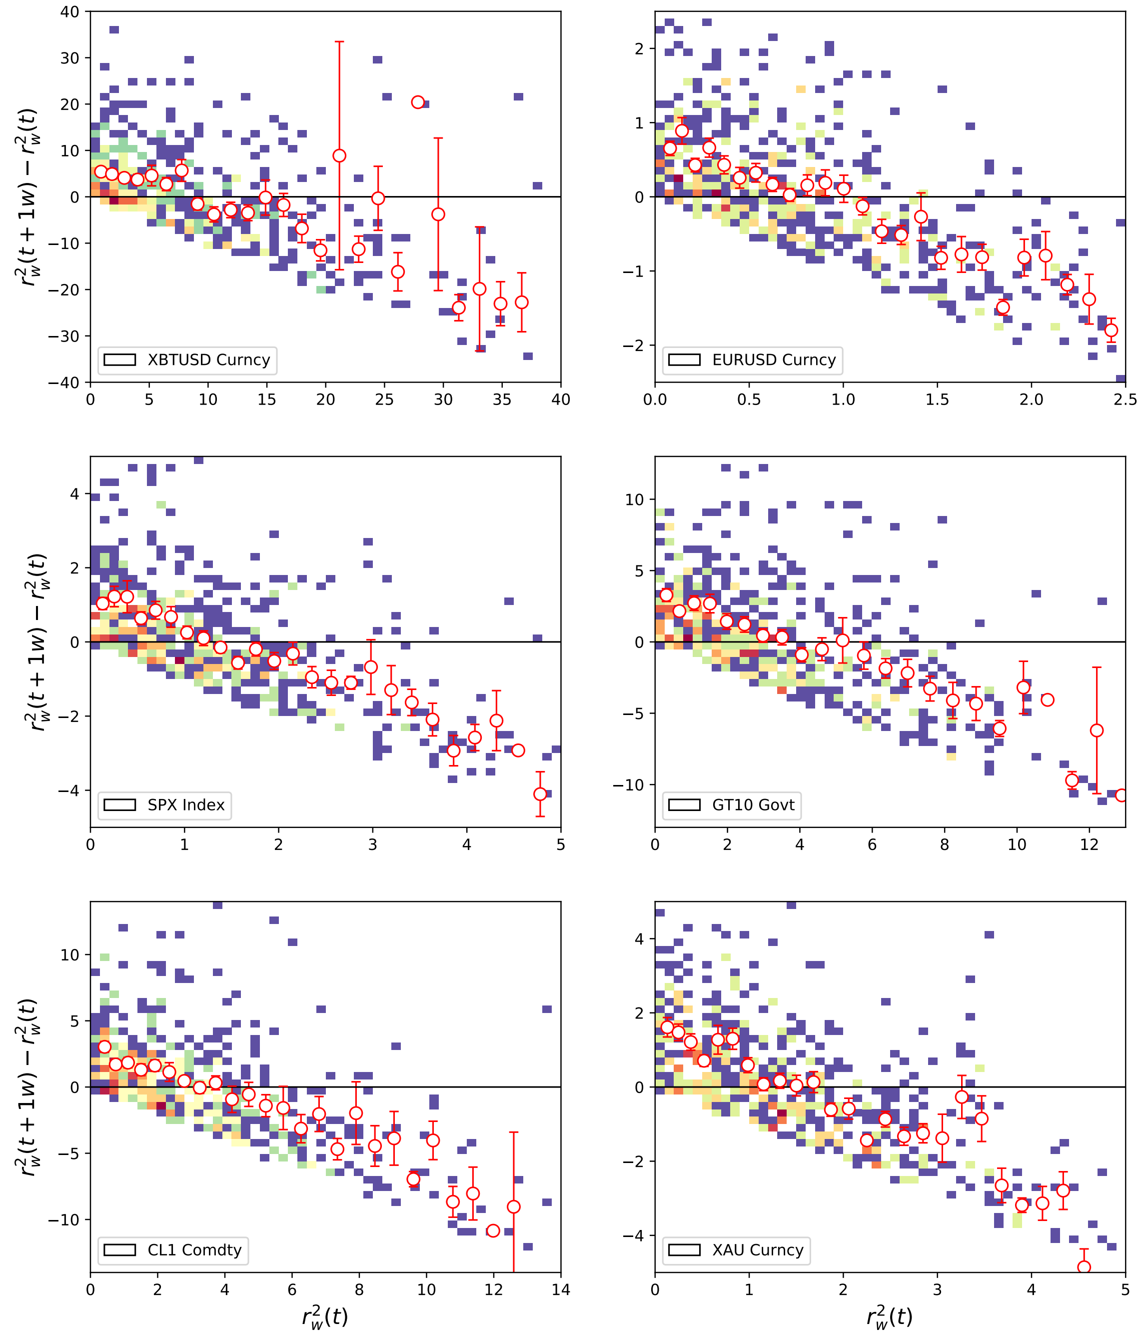

Figure 2: Increments of squared weekly returns against starting values. Circles and error bars are means and their standard deviations of y-values in x-bins. Data range: 2011-01 to 2021-12.

2 Volatility Predictability

2.1 Model Performance Metrics

In this study we use the Pearson correlation coefficient as a model performance metric, with the focus on measuring how closely predicted volatilities follow the *trend* of observed, disregarding the absolute scale. The Root Mean Squared Errors (RMSE) is also provided for out-of-sample performance, as a complementary metric to measure *scale*. It serves as a check on the correlation metric.

The correlation coefficient can be a better metric than RMSE for circumstances where matching the trend is predominately important. To illustrate this, let's consider the following scenario:

Observed/Predicted	Values
Observed	$y_1, y_2, \cdots, y_{2n}$
Predicted by Model 1	$y_1 + \Delta, y_2 + \Delta, \cdots, y_{2n-1} + \Delta, y_{2n} + \Delta$
Predicted by Model 2	$y_1 + \Delta, y_2 - \Delta, \cdots, y_{2n-1} + \Delta, y_{2n} - \Delta$

where $2n$ is the number of observations, the predicted values by model 1 is the observed plus a constant Δ, and the predicted values by model 2 is the observed plus Δ for odd indices and minus Δ for even indices. The RMSE for both models is Δ. The correlation coefficient is 1 for model 1 and smaller than 1 for model 2. While model-1 predictions follow precisely the trend of observed, model-2 predictions oscillate about the observed. In this case, RMSE misses the point, but correlation coefficient captures it.

On another note, the correlation coefficient tends to be less sensitive to outliers than RMSE: outliers appear in both the numerator and denominator of its formula, reducing their overall effects. This is important considering that jumps are present in volatility timeseries (Bates 1996, Barndroff-Nielsen and Shephard 2004, Todorov and Tauchen 2011).

2.2 HAR Model

Corsi (2009) designed a heterogeneous autoregressive (HAR) model to parsimoniously capture the strong persistence typically observed in realized variance (RV). The idea was motivated by the so-called Heterogeneous Market Hypothesis (Müller et al. 1993), and the model put the focus on the heterogeneity that originates from the difference in the temporal horizons. This temporal heterogeneity was considered a result of financial market participants having a large spectrum of trading frequencies. The traditional ARCH model (Engle 1982) is generally considered too simplistic to satisfactorily describe the long-run dependencies in most realized volatility series (see, e.g., Bollerslev et al. 2016).

The HAR model specifies RV as a linear function of daily, weekly and monthly realized variance components, and can be expressed as

$$\mathrm{RV}_{t+1d}^{(d)} = c + \beta^{(d)}\mathrm{RV}_t^{(d)} + \beta^{(w)}\mathrm{RV}_t^{(w)} + \beta^{(m)}\mathrm{RV}_t^{(m)} + \omega_{t+1d} \tag{2.1}$$

where ω_{t+1d} is an error term, c, $\beta^{(d)}$, $\beta^{(w)}$, and $\beta^{(m)}$ are model parameters, $\mathrm{RV}_{t+1d}^{(d)}$ is the realized variance of day $t+1d$, and $\mathrm{RV}_t^{(d)}$, $\mathrm{RV}_t^{(w)} = \sum_{i=1}^{wd} \mathrm{RV}_{t-i+1}^{(d)}/wd$, and $\mathrm{RV}_t^{(m)} = \sum_{i=1}^{md} \mathrm{RV}_{t-i+1}^{(d)}/md$ denote the daily, weekly and monthly lagged realized variance as of day t, respectively, with wd denoting the number of days in a week and md denoting the number of days in a month. In this study, we set $wd = 7$ and $md = 30$ for cryptocurrencies and $wd = 5$ and $md = 22$ for other instruments.

To ensure the positiveness of estimated $RV_{t+1d}^{(d)}$, we use logarithm-scaled RV in Equation 2.1 in this study. We also use the same model for h-day ahead prediction:

$$\log(\mathrm{RV}_{t+h}^{(d)}) = c_h + \beta_h^{(d)} \log(\mathrm{RV}_t^{(d)}) + \beta_h^{(w)} \log(\mathrm{RV}_t^{(w)}) + \beta_h^{(m)} \log(\mathrm{RV}_t^{(m)}) + \omega_{t+h} \tag{2.2}$$

where the set of model parameters are specific to h.

2.2.1 In-Sample Model Performance

Figure 3 shows the HAR model fitting results for a range of prediction horizons and the instruments in different asset classes. The correlation coefficients between observed and predicted decrease with increasing prediction horizon across the board. Among the 10 cryptocurrencies, XBT shows consistently higher correlations than for other coins. In general, the fixed income instruments show higher correlations than for other instruments, with lowest correlations observed for the traditional currencies. The similarity in the Correlation-Horizon curves among the asset classes illustrates that the crypto is not exceptional in volatility predictability.

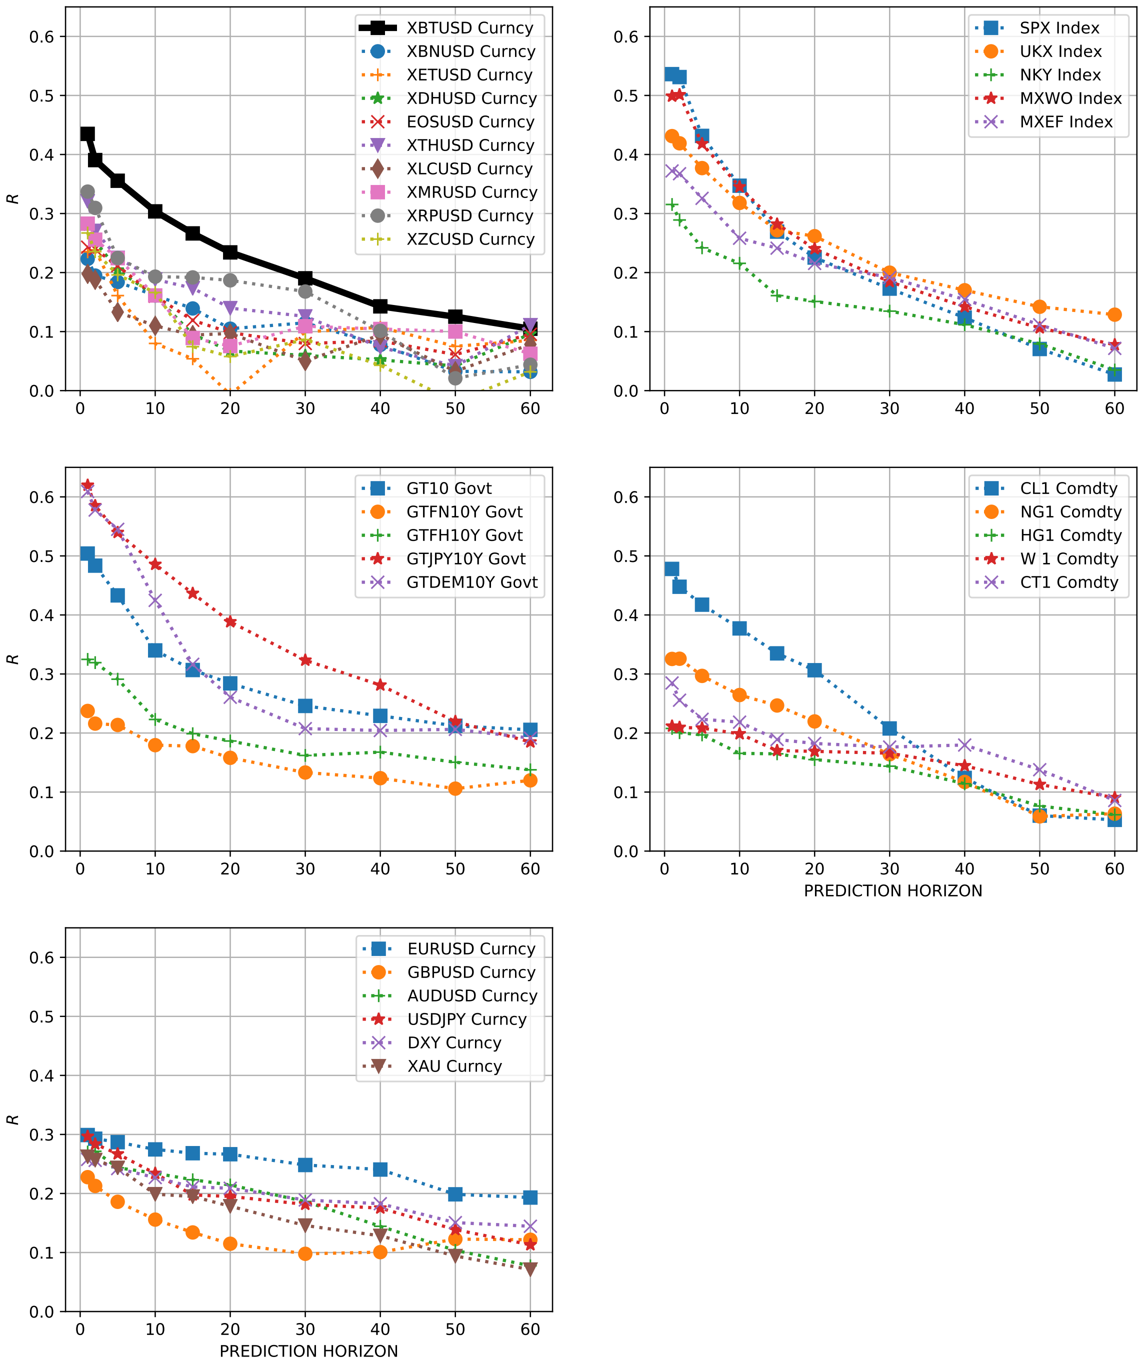

Figure 3: In-sample correlation coefficients (R) between observed and predicted *volatilities* for cryptocurrencies, equity indices, government and agency bonds, commodities, and currencies. The absolute return is used as a proxy for observed volatility. $exp(./2)$ is applied to output of Equation 2.2 for predicted volatility. Data range: 2011-01 to 2021-12.

2.2.2 Out-Of-Sample Model Performance

Figure 4 shows the HAR model out-of-sample prediction results for a range of prediction horizons and the instruments in different asset classes. In general, the correlation coefficients between observed and predicted decrease with increasing prediction horizon, with a slight tipping up at long horizons (≥ 40) for the crypto and equity index categories. Among the 10 cryptocurrencies, XBT shows higher correlations than for other coins at most horizons. The fixed income instruments show higher correlations than for other instruments. For XBT, its correlations are within the range for other asset classes. This again illustrates that the crypto is not exceptional in volatility predictability.

2.3 Random Forest Method

Although volatility demonstrates strong persistence, which justifies models such as HAR model, the model performance as measured by the correlation coefficient is not quite satisfactory. We explore Random Forests (RF) (Breiman 2001) to see if we can improve prediction accuracy by drawing information from other instruments and by relaxing the linearity assumption.

For a given target instrument, all the 35 instruments listed in Table 1, plus the ones in *Others*, are used as RF model input features. To include realized variance with past memory longer than one month, we create a quantity denoted as RV^e for the exponentially weighted moving average (EWMA): $\mathrm{RV}_t^e = \lambda \mathrm{RV}_{t-1}^e + (1-\lambda) r_t^2$, where the decay factor λ is set to 0.97. This value corresponds to a half-life of 23 days. Specifically, the input training dataset of (feature vector, label) pairs is compiled as follows:

$$
\begin{aligned}
X_t = &\Big[\Big(\mathrm{RV}(i)_t^d, \mathrm{RV}(i)_t^w, \mathrm{RV}(i)_t^m, \mathrm{RV}(i)_t^e, \mathrm{SIGN}(r_t(i)) | i = 1, 2, \cdots, 35\Big), \\
&\mathrm{VX}_t^d, \mathrm{VX}_t^w, \mathrm{VX}_t^m, \mathrm{VX}_t^e, \mathrm{M2}_t^d, \mathrm{M2}_t^w, \mathrm{M2}_t^m\Big] \\
y_t = &|r_{t+h}|
\end{aligned}
$$

where t is the index of sample dates, i is the index of instruments, h is the prediction horizon, r_{t+h} is the h-ahead return of the target instrument, VX^d, VX^w, VX^m and VX^e are respectively daily, weekly, monthly, and EWMA average of squared VIX values, and $\mathrm{M2}^d$, $\mathrm{M2}^w$ and $\mathrm{M2}^m$ are respectively daily, weekly and monthly average of M2 Index values.

At time indexed by t, we predict h-ahead volatility using the model trained with rolling latest dataset of size m: $\{X_s, y_s | s = t-h-m+1, \cdots, t-h\}$. $m = 120$ is used for this study, which is about the number of past observations with weights greater than 0.001 for volatility estimation using EWMA with the recommended decay factor of 0.94 (J.P.Morgan/Reuters 1996).

A justification for treating instruments as *features* is the expectation that they represent different dimensions or aspects of the complex market: financial, economic, monetary, etc., since the instruments are selected from different asset classes and carry over diverse information.

Figure 5 shows the out-of-sample correlation coefficient comparison between the RF and HAR models for six representative instruments from major asset classes. It shows that RF outperforms HAR with only a few exceptions (For XBTUSD, HAR outperforms RF by small margins for prediction horizons 40 and 50). Figure 6 compares the out-of-sample Root Mean Squared Errors (RMSE) of the two models. For XBTUSD, the result is mixed: RF has smaller RMSE for prediction horizon < 10, but the pattern is reversed for horizon ≥ 15. For the other 5 instruments, RMSE of RF is mostly smaller than that of HAR, and for certain prediction horizons, it's much smaller.

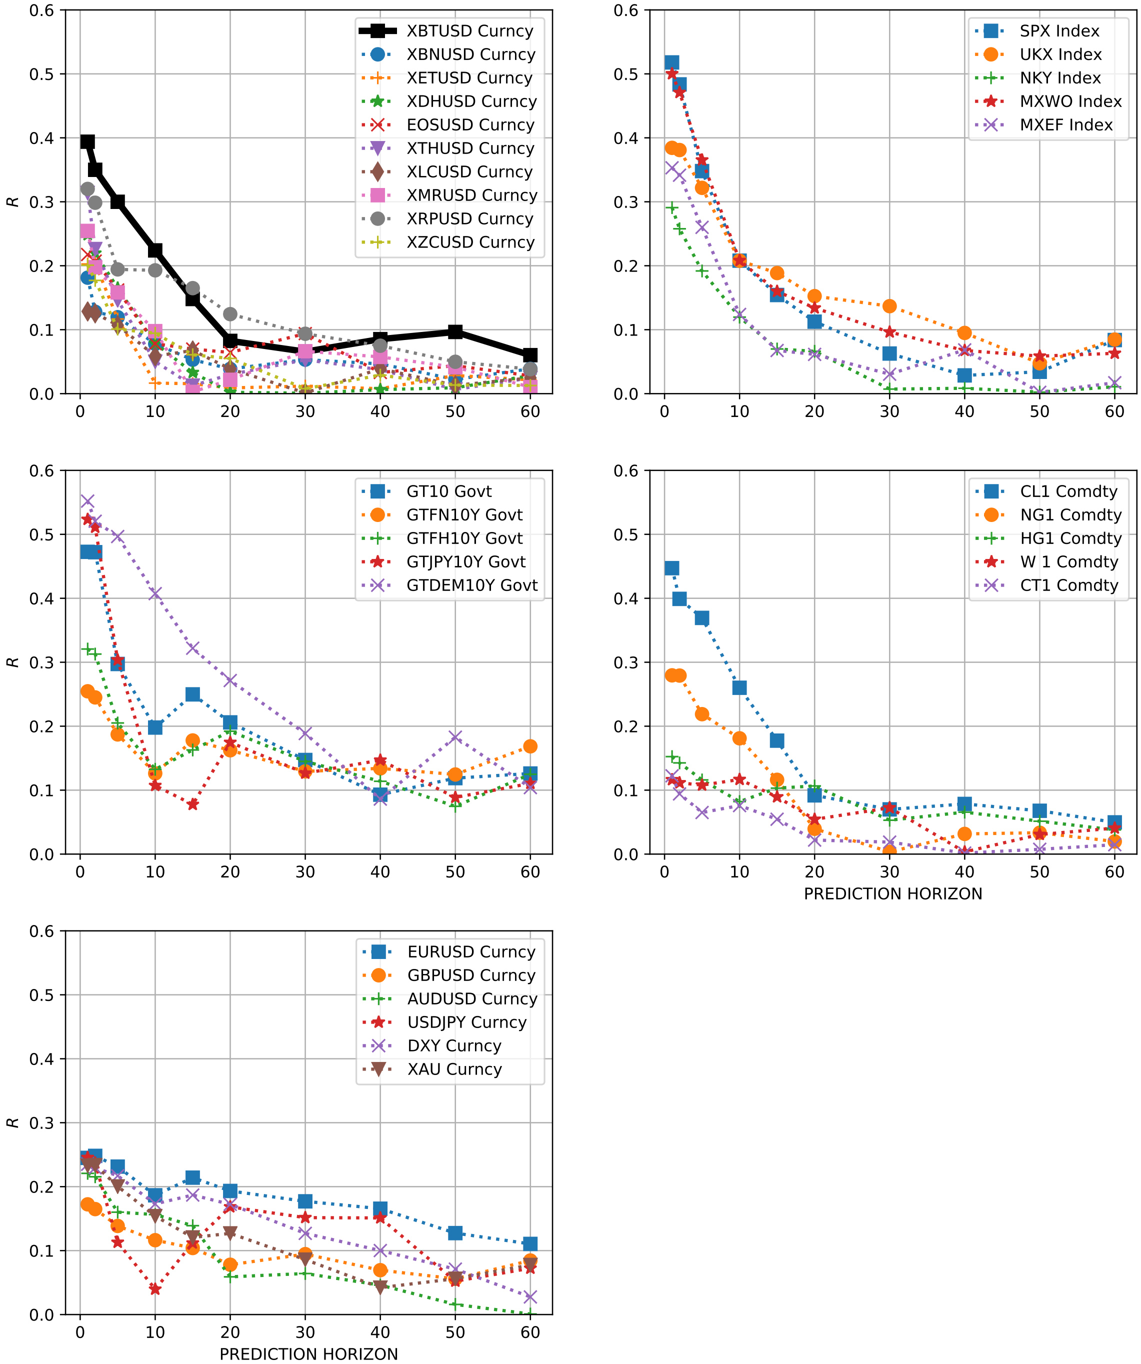

Figure 4: Out-of-sample correlation coefficients (R) between observed and predicted *volatilities* for cryptocurrencies, equity indices, government and agency bonds, commodities, and currencies. The absolute return is used as a proxy for observed volatility. $exp(./2)$ is applied to output of Equation 2.2 for predicted volatility. Each prediction is made using a model fitted with a rolling window of 120 days. Data range: 2011-01 to 2021-12.

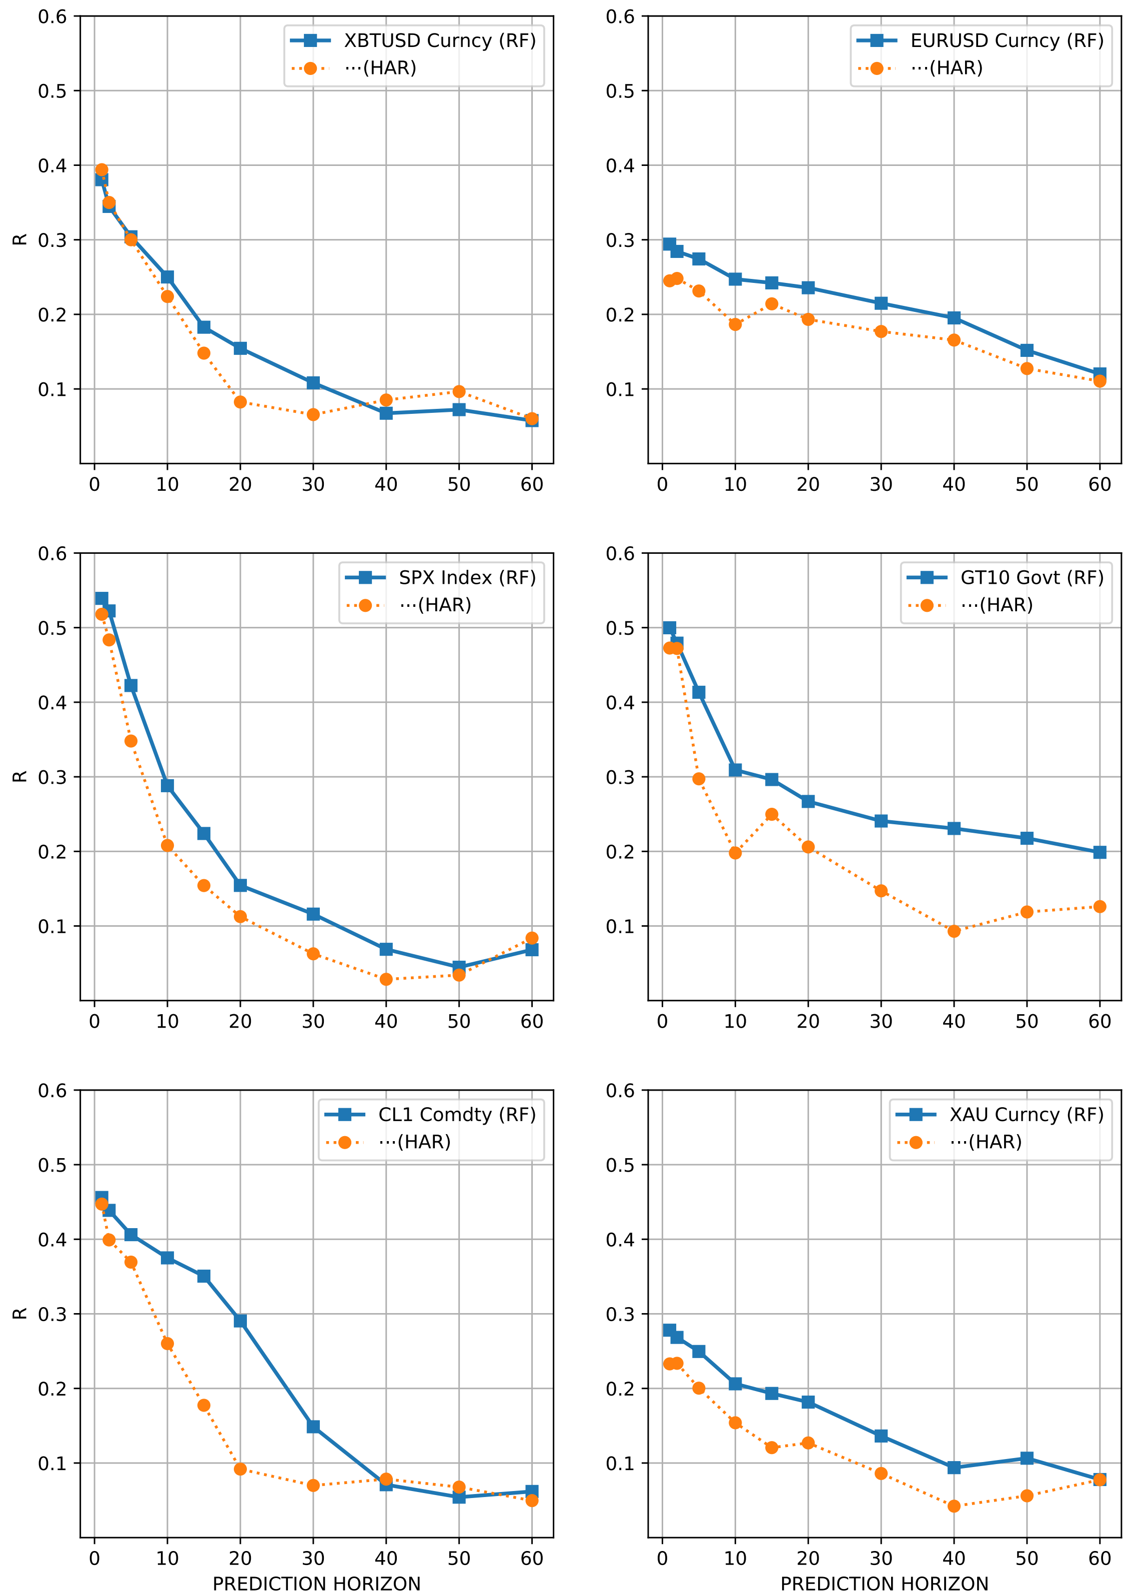

Figure 5: Comparing out-of-sample correlation coefficients (R) between observed and predicted *volatilities*: points with square marks correspond to RF prediction results; those with circle marks correspond to HAR results. The six instruments represent major asset classes. The absolute return is used as a proxy for observed volatility. The scikit-learn RandomForestRegressor is used as the RF model implementation, and the training parameters are: $n\_estimators = 300$, $max\_depth = 3$, and $max\_features = 10$. Data range: 2011-01 to 2021-12.

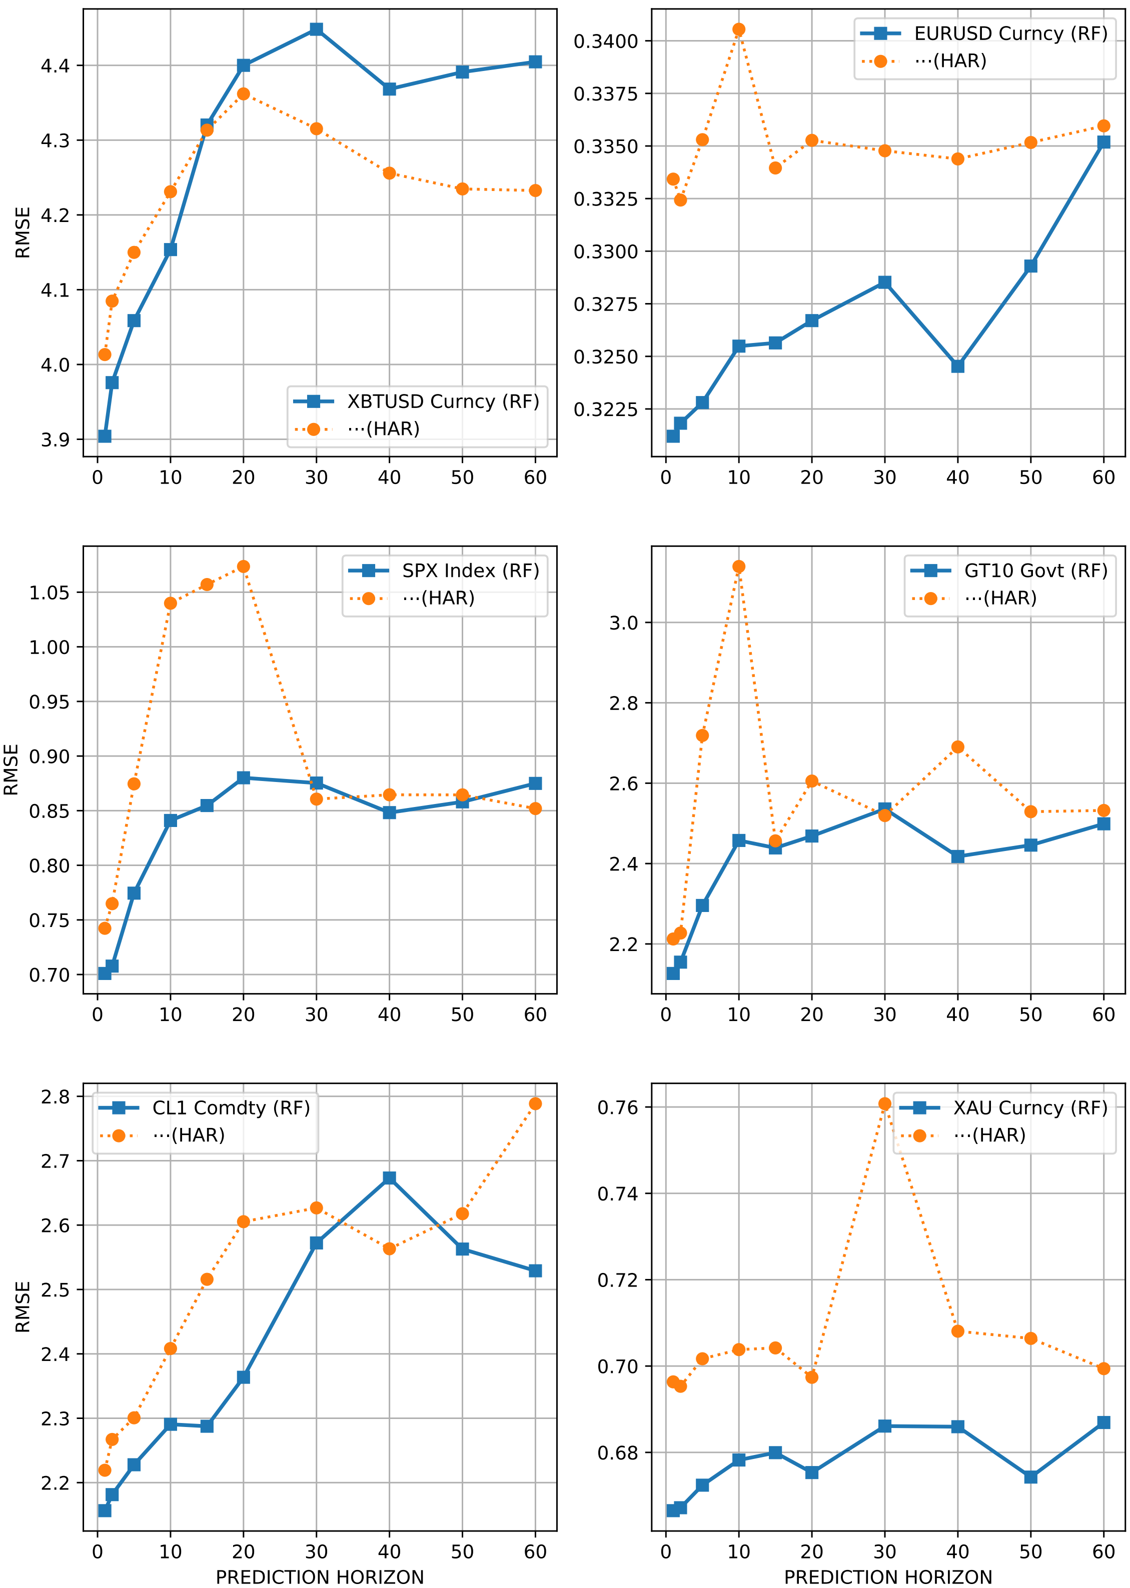

Figure 6: Comparing out-of-sample model RMSEs: points with square marks correspond to RF prediction results; those with circle marks correspond to HAR results. The six instruments represent major asset classes. The absolute return is used as a proxy for observed volatility. The scikit-learn RandomForestRegressor is used as the RF model implementation, and the training parameters are: $n\_estimators = 300$, $max\_depth = 3$, and $max\_features = 10$. Data range: 2011-01 to 2021-12.

3 Discussion

Figure 7 shows average feature importance for different prediction horizons. Note the following:

- for horizon 1, auto-feature RV^d ("XBTUSD Curncy" in the figure) is on the top, whereas for horizons 5, 10 and 20, auto features fall off the top five.
- for large horizons 40 and 60, auto-features RV^e ("XBTUSD Curncy(EWMA)" in the figure) come back on top three.
- except for horizon 1, the EWMA or monthly lagged features dominate.

This figure suggests that the persistence factor dominates predictions for small horizons ($H < 5$), diminishes with increasing horizon up to an intermediate level of 20, then comes back at large horizons (≥ 40). This explains the R-HORIZON shapes for the HAR model (see the first subplot in Figure 4). On the other hand, the RF model performs better for intermediate horizons by utilizing information from other instruments to make up for the diminishing persistence factor.

We have also explored other machine learning models for the volatility prediction. One such model is LSTM (Hochreiter and Schmidhuber 1997), which is capable of dealing with timeseries input. Unfortunately, no results better than with RF have been obtained. However, we don't conclude that the LSTM model is not suitable here since we haven't performed tuning of hyper parameters thoroughly. We leave further exploration of machine learning techniques on volatility prediction to future work.

From the correlation-prediction horizon curves (Figure 3 and Figure 4), we see that the crypto asset class behaves in line with many instruments in other asset classes. We conclude that the volatility of cryptocurrencies is on average as predictable as for other instruments.

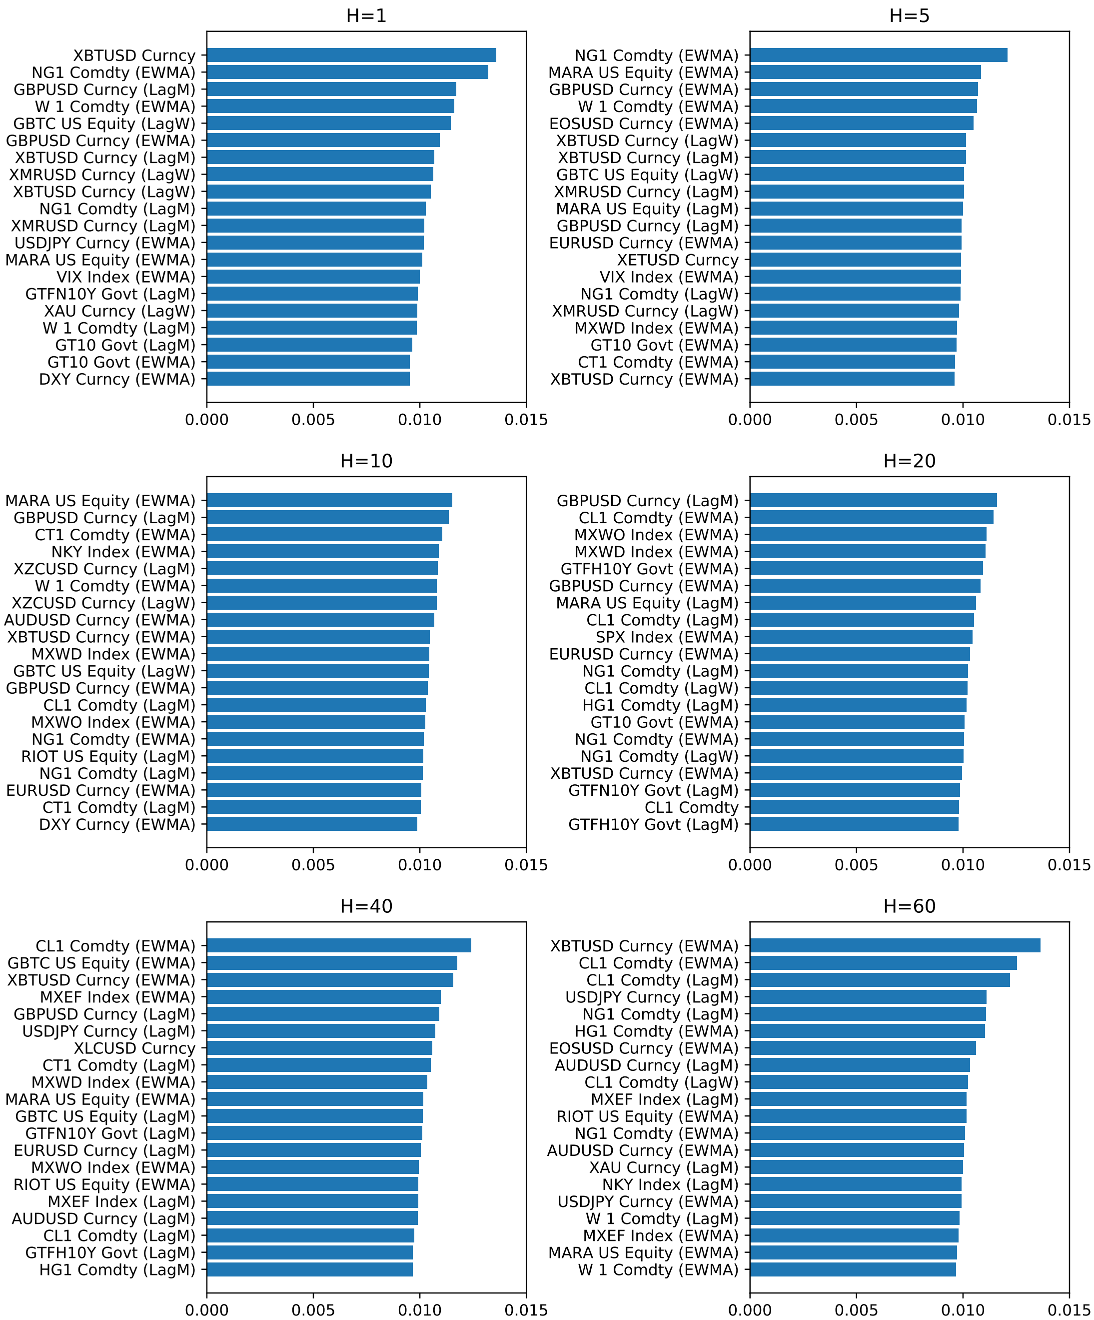

Figure 7: Average feature importance of RF models trained with rolling window of 120 days. H: prediction horizon. Data range: 2011-01 to 2021-12.

References

Barndroff-Nielsen, O. and N. Shephard (2004). Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometric 2*(1), 1–37.

Bates, D. S. (1996). Jumps and stochastic volatility: Exchange rate processes implicity in deutsche mark options. *The Review of Financial Studies 9*(1), 69–107.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics 31*(3), 307–327.

Bollerslev, T., A. Patton, and R. Quaedvlieg (2016). Exploiting the errors: A simple approach for improved volatility forecasting. *Journal of Econometrics 192*(1), 1–18.

Box, G., G. Jenkins, and G. Reinsel (2008). *Time Series Analysis, Forecasting and Control (4th ed.)*. John Wiley & Sons, Ltd.

Breiman, L. (2001). Random forests. *Machine Learning 45*(1), 5–32.

Cont, R. (2007). Volatility clustering in financial markets: Empirical facts and agent-based models. In *Long Memory in Economics*, pp. 289309. Springer.

Cont, R., J.-P. Bouchaud, and M. Potters (1997). Scaling in financial data: Stable laws and beyond. In *Scale Invariance and Beyond*, Volume 52. Springer.

Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics 7*(2), 174–196.

Cox, J. C., J. E. J. Ingersoll, and S. A. Ross (1985). A theory of the term structure of interest rates. *Econometrica 53*(2), 385–407.

CryptoFundResearch (2021). 2021 q3 crypto fund quarterly report. Technical report. https://cryptofundresearch.com/q3-2021-crypto-fund-report/.

Ding, Z., C. W. Granger, and R. F. Engle (1993). A long memory property of stock market returns and a new model. *Journal of Empirical Finance 1*(1), 83–106.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica 50*(4), 987–1007.

Guillaume, D., M. Dacorogna, R. Dave, U. Müller, R. Olsen, , and O. Pictet (1997). From the birds eye view to the microscope: A survey of new stylized facts of the intraday foreign exchange markets. *Finance and Stochastics 1*, 95131.

Heston, S. L. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies 6*(2), 327–343.

Hochreiter, S. and J. Schmidhuber (1997). Long short-term memory. *Neural computation 9*(8), 1735–1780.

J.P.Morgan/Reuters (1996). Riskmetricstechnical document. Technical report, J.P.Morgan/Reuters.

Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business XXXVI*, 392417.

Müller, U. A., M. M. Dacorogna, R. D. Dave, O. V. Pictet, R. B. Olsen, and J. R. Ward (1993, oct). Fractals and intrinsic time a challenge to econometricians. In *39th International AEA Conference on Real Time Econometrics*.

Nelson, D. (1990). Arch models as diffusion approximations. *Journal of Econometrics 45*, 7–38.

Papanicolaou, G. and K. Sircar (1999, 02). Mean-reverting stochastic volatility. *International Journal of Theoretical and Applied Finance 3*.

PwC (2021). 3rd annual global crypto hedge fund report 2021. Technical report. https://www.pwc.com/gx/en/financial-services/pdf/3rd-annual-pwc-elwood-aima-crypto-hedge-fund-report-(may-2021).pdf.

Reiff, N. (2022). Why is bitcoin volatile? *Investopedia*. https://www.investopedia.com/articles/investing/052014/why-bitcoins-value-so-volatile.asp.

Sigalos, M. (2021). Bitcoin's wild price moves stem from its design you'll need strong nerves to trade it. *CNBC*. https://www.cnbc.com/2021/05/19/why-is-bitcoin-so-volatile.html.

Todorov, V. and G. Tauchen (2011). Volatility jumps. *Journal of Business & Economic Statistics 29*(3), 356–371.