Pure Momentum in Cryptocurrency Markets

By CESARE FRACASSI AND SHIMON KOGAN*

Momentum is one of the most widespread, persistent, and puzzling phenomenon in asset pricing. The prevailing explanation for momentum is that investors under-react to new information, and thus asset prices tend to drift over time. We use a unique feature of cryptocurrency markets: the fact that they are open 24/7, and report returns over the last 24 hours. Thus, the one-day return is subject to predictable fluctuations based on the removal of lagged information. We show that investors respond positively to changes in reported returns that are unrelated to any new release of information, or change in the asset fundamentals. We call this behavioral anomaly "Pure Momentum".

One of the most well studied asset pricing patterns is the positive relationship between an asset's returns and its lagged price performance (e.g. Jegadeesh and Titman (1993), Asness, Moskowitz and Pederson (2013)), called "momentum". This phenomenon appears to hold across many different asset classes (individual stocks, international stock markets, government bonds, currencies, and commodities) and thus has been at the center of the anomalies debate. In fact, Fama admitted that momentum was the biggest embarrassment to the theory of financial market efficiency.

As puzzling as momentum may appear, there is no definitive explanation for it. Some papers propose risk-related explanations (e.g., Daniel and Moskowitz (2016)) but rationalizing the high Sharpe Ratio on a dynamic momentum strategy is difficult with reasonable risk aversion parameters. Other papers focus on behavioral explanations that center on under-reaction to information (e.g., Daniel, Hirshleifer and Subrahmanyam (1998)) but it is difficult to provide direct evidence for this channel. Finally, investors may respond to past returns themselves if they have extrapolation bias or diagnostic beliefs (e.g., Bordalo et al. (2019)). Indeed, it is difficult to isolate the sheer effect of past returns from other correlates – after all, investors may react to some latent factors that, at the same time, affected prices too. In short, isolating the pure response to past prices, which we call

* Fracassi: Coinbase Institute, and University of Texas at Austin. Kogan: Reichman Univesity and the Wharton School. The views and opinions expressed herein are those of the authors and do not necessarily reflect the views of Coinbase or its employees and summarizes information and articles with respect to cryptocurrencies or related topics that the author believes may be of interest. This material is for informational purposes only, and is not (i) an offer, or solicitation of an offer, to invest in, or to buy or sell, any interests or shares, or to participate in any investment or trading strategy, (ii) intended to provide accounting, legal, or tax advice, or investment recommendations or (iii) an official statement of Coinbase. No representation or warranty is made, expressed or implied, with respect to the accuracy or completeness of the information or to the future performance of any digital asset, financial instrument or other market or economic measure. The information is believed to be current as of the date indicated on the materials. Recipients should consult their advisors before making any investment decision.

"pure momentum" is empirically very challenging.

In this paper, we identify a unique natural experiment that naturally occurs in cryptocurrency markets where the perception of past returns is shocked without any new information or change to fundamentals. In traditional securities markets, returns are usually defined relative to a fixed reference point, usually close price of last trading day. However, In cryptocurrency markets, trading takes place continuously 24/7 around the world. Reporting returns requires a reference point, and thus the most salient return windows is the last 24 hours return (see, for example, Figure 1). A key feature of these returns is that the reference point for computing return is not fixed, but moves continuously over time as the return window constantly shifts. Thus, the one-day return is subject to predictable fluctuations based on the removal of lagged changes in stock price. To illustrate, imagine that the price of an asset was \$105 24 hours ago, \$100 23 hours ago, and back to \$105 currently. While the current 24-hour return is zero, we can anticipate daily return next hour to be positive 5%, other things equal, since the reference price used to calculate the 24-hour return will be lower next hour. That is, the reference price for the return calculation moves, in this example, from \$105 to \$100. While such a change does not reveal new information – after all, the price 23 hours ago is well know – it may affect the perceived value of the asset if investors rely on the one-day returns to update their beliefs. It is important to note that this feature is unique relative to many other markets. For example, foreign exchange markets, which do trade around the clock as well, typically display day returns relative to a fixed reference point, which is 12AM UTC.

We utilize this predictive shift in reference points to evaluate how predictable changes in returns affect future returns. Indeed, we find strong evidence to support the presence of pure momentum. Using high-frequency trading data on 138 cryptocurrencies over a period of 5.5 years, we show that the current 1-hour returns is negatively related to lagged 1-hour returns exactly 24 hours ago: As positive lagged-24 returns disappear from the reference window, returns appear to decline, inducing investors to sell in response to the change in return, a pure momentum effect.

This effect is not present before the 24 hour mark and dissipates within 4 hours. Likewise, the 15-min returns responds strongly to the 15-min returns 24 hour prior, but not to the 15-min returns just prior to the 24 hour mark. The results appear to be robust over time, i.e., across years, months of the year, days of the week, and time of day. In the cross-section of cryptocurrencies, the effect appears to be more pronounced for more liquid currencies, suggesting that it is not a "small-cap effect", but it is due to an "excess of attention" on the part of investors. Finally, the economic magnitude of the effect is very large. A long-short strategy that goes long (short) the two cryptocurrencies with the lowest (highest) 24-lagged 1-hour returns, earns an annualized return of 538%, and a Sharpe ratio of 3.52, before trading costs. Given the high turnover in this strategy, it becomes unprofitable after taking into account the typical trading fees and the trading price impact.

Our results are most closely related to Phillips, Pukthuanthong and Rau (2016), who study the response of mutual fund flows to fixed-length, regulatory mandated holding period returns. They find that mutual fund flows respond to predictable changes in holding period returns, and that mutual fund manager time their advertising to coincide with these changes. While their instrument is similar to ours, and their results are consistent with our findings, there are substantial difference in the markets that we study. Specifically, it is much harder to arbitrage inefficient open-ended mutual fund flows than inefficient price formation in what are some of the must heavily traded markets. In the empirical asset pricing literature, periodicity in returns around salient intervals have been documented. Jegadeesh (1990) shows that monthly returns are correlated with the same month returns one, two and three years prior. Heston and Sadka (2008) extend the analysis by showing that the same patterns extends as far back as twenty years. Both papers provide strong empirical evidence but do not suggest a mechanism. Our results suggest that, to the extent that there investors measure return around these salient intervals, these period return pattern can result from pure momentum.

Traditional momentum strategies have also been studied in cryptocurrency markets, with inconclusive results: Grobys and Sapkota (2019) find no momentum in sample 2014-2018. Shen, Urquhart and Wang (2020) use a three factor model, and find a reversal factor, a negative momentum for the period 2013-2019. Jia, Goodell and Shen (2022) uses a more recent sample period and find positive momentum. All these studies use short time periods to estimate a factor model, and thus very dependent on the chosen sample period. Our paper uses a unique reporting feature of cryptocurrency markets to predict the investment behavior of traders. The identification strategy is high frequency, and the effect is persistent across different sample periods.

I. Pure Momentum

Currently U.S. stock markets are open from 9.30am to 4pm ET. The rationale for limited trading hours is that concentrated liquidity helps market stability and efficiency. Furthermore, most company announcements occur during off-trading hours, to avoid insider trading.

The amount of retail trading has increased significantly since the beginning of the COVID pandemic. Estimates from NASDAQ show trading volume doubling right after the March 2020 pandemic lockdowns, and remaining at high level ever since.[1] With such a rise in retail trade share, some started advocating for stock markets to remain open for longer hours to cater to the increased demand for trading. Many point to cryptocurrency trading venues as examples of markets that are open 24/7, 365 days a year.

[1]NASDAQ - A New Way to Look at Retail Trading Trends. https://www.nasdaq.com/articles/a-new-way-to-look-at-retail-trading-trends

An aspect to be considered when markets are open 24/7 is how to report past returns. In traditional securities markets, the close of the previous trading day is used as reference point to compute intra-day and daily returns. But crypto and Forex markets do not close, and thus do not have a natural reference price. One solution is to choose an arbitrary reference time. For example, Forex markets reset returns at 12am UTC (Coordinated Universal Time). Crypto-markets for the most part instead report returns using a rolling reference window of 24 hours. Figure 1 shows how Coinbase reports 24-hour returns in its platform. Table 1 shows the reference price convention for some of the top crypto exchanges and brokers. All crypto-exchanges use a 24-hour rolling window and some of the brokers use a 12AM reference point when reporting daily returns.

While no reporting is intrinsically better than another, the 24-hour rolling window is more prone to possible misinterpretation by unsophisticated investors. Innovations in returns are driven by two components, one is the change in current (hourly) price from t to $t+1$, and the second is the fact that the reference window moves, and thus older returns are dropped from the window.

More formally, we define $ret(t)$ as the return over the last 24 hours:

$$\text{ret(t)} = \frac{P_t}{P_{t-24}} - 1$$

One hour later, the 24-hour return is:

$$\text{ret(t+1)} = \frac{P_{t+1}}{P_{t-23}} - 1$$

The change in return from t to $t+1$ is thus:

$$\Delta\text{ret} = \text{ret}(t+1) - \text{ret}(t) = \frac{P_{t+1}}{P_{t-23}} - \frac{P_t}{P_{t-24}} = \left(\frac{P_t}{P_{t-23}}\right)\left(\frac{P_{t+1}}{P_t} - \frac{P_{t-23}}{P_{t-24}}\right)$$

From the equation above, we can see that the change in return is driven by unexpected innovation in prices $\left(\frac{P_{t+1}}{P_t}\right)$ and by the removal of stale returns from the reference window $\left(\frac{P_{t-23}}{P_{t-24}}\right)$.

Our identification strategy relies on the assumption that unsophisticated investors might notice returns changing, and not realize that it might be driven by the removal of stale information. If this happens, we have a clean exogenous shock to returns that is not influenced by other confounding omitted variable that might simultaneously influence innovations in returns. Using such instrument, we can measure how investors respond to pure changes in returns that are unaffected by fundamentals, a “pure momentum”.

II. Data and Experiment Design

Our data consistent of the complete set of cryptocurrencies listed to trade on Coinbase, the largest cryptocurrency exchange in the US, traded between July 1st, 2015 to December 31st, 2021. We collected hourly and 15-minutes candles using Coinbase's public APIs.[2] For each currency, we calculate the 1-hour and 15-min returns. Table 2 presents the summary statistics of our data. While the sample period starts in 2015, the number of cryptocurrencies, and thus the number of observations, during the first 4 years of our sample period is rather limited. Thus, our analysis implicitly loads more heavily on the last three years of data. Not surprisingly, our sample is characterized by large, and volatile, average returns. The average hourly return in the sample is 1.5bps, with a hourly standard deviation of 1.66%. Hourly trading volume average $\$1.6M$.

The main analysis is straightforward: we estimate the incremental relation between time t returns and each of the lagged 1-hour interval returns, by regressing time t hourly return for crypto i on all its lagged 1-hour returns spanning 60 lags.

Specifically, we run the following OLS regression:

$$r_{(c,t\to t+1)} = \alpha_0 + \sum_{k=1}^{60} \alpha_k r_{(c,t-k\to t-k+1)} + \gamma_{hour} + \gamma_{day} + \gamma_{month} + \gamma_{year} + \gamma_c \tag{1}$$

where $ret_{(c,t\to t+\delta)}$ is the return for cryptocurrency c from hour t to hour $t+\delta$, γ_{hour} are hour-of-the-day dummies, γ_{day} are day-of-the-week dummies, γ_{month} are month-of-the-year dummies, γ_{year} are year dummies, and γ_c are cryptocurrency FE. Standard errors are clustered by time.

III. Main Results

The results of the OLS estimation of equation 1 are displayed in Figure 2. The mark represents the estimate of the coefficient, and the line the 95% confidence interval. The results are striking – while most lagged returns have no statistically significant relations with subsequent returns, we find strong and negative relation between time t returns and lagged 1-hour returns 24 hours earlier. Notice that lagged 23 and 22 hour returns are not informative for future returns.[3] The negative serial correlation is also economically significant. A 1% increase in prices from $t-24$ to $t-23$ leads to a 3.4bps decrease in returns form time t to time $t+1$.

What can explain this patter? The reporting of cryptocurrency returns with shifting anchor are consistent with our findings. When the 1-hour return 24 hours ago is negative (positive), its removal through the shift by one hour will increase

[2]https://docs.cloud.coinbase.com/sign-in-with-coinbase/docs/api-users

[3]We also find strong significant negative correlation between time t and time $t-1$ and $t-2$ returns, which are most likely driven by micro-structure market frictions, and are not part of our study.

perceived returns measured over a 24-hour window, causing investors affected by extrapolation biases to buy the cryptocurrency in the hope of continuously future price appreciation, thus increasing price pressure and higher returns.

An important feature of this natural experiment is that investors have to frequently pay attention to the return, in order for them to be affected by the changing returns driven by the disappearance of stale information. In other words, opposite to the "limited attention" literature, this specific behavioral bias is triggered by an "excess attention" to crypto returns. In fact, the effect appears to persists for 4 hours, consistent with the idea that traders respond quickly to this information when they frequently check crypto-returns.

It is also interesting to note that the negative loading on lags 24 (and 25) is repeated in lags 48 (and 49), while the magnitude is muted. Given that 2-day returns are not saliently reported, this results in consistent with pure momentum being echoed in prices through the endogenous response of traders to the initial shock.

To sharpen the identification, we repeat the analysis with 15-min returns (increasing the number of lags to 120, spanning the previous 30 hours). Figure 3 shows that the 15-min return exactly 24 hours prior (but not 23 hours or even 23.75 hours) is strongly negatively related to subsequent 15-min returns.

IV. Cross-sectional Tests

To assess the robustness of the results, and to validate the channels through which past stale returns influence future returns, we study how the 24-hour pure momentum effect varies across time and cryptos. In all these tests, we compare the lag-24 hour coefficient with the lag-22 hour coefficient, the latter used as a sort of placebo test.

First, we estimate the lag 24-hour coefficient separately for different years, months of the year, days of the week, and times of the day. Prior literature found seasonal patterns in momentum. For example, Sias (2007) shows that momentum most pronounced during year-end months. Figure 4 shows the coefficient estimated for the different years in our sample period. Contrary to the view that our effect is driven by investors in under-developed markets, we find the coefficient becoming statistically significant, and similar in magnitude, for the last 4 years of our sample. As investors pay more and more attention to crypto returns, they are more subject to this bias. In contrasts, lag-22 return coefficients are all statistically insignificant, and do not show any trends over time.

Second, we proceed to test whether cryptocurrencies that are more frequently traded are more or less affected by this pure momentum effect. Each hour, we split the sample into terciles, and then run equation 1 for each tercile. Figure 5 shows that the lag-24 hour coefficient becomes *more* negative when volume, and thus liquidity, is high. This is consistent with the evidence that investors' attention and cryptocurrency liquidity are related, making them more subject to this extrapolation bias. We do not find any patters or statistically significant effect for the placebo lag-22 return coefficient estimates. Third, we look at whether the effect is linear in the magnitude of the lagged returns or symmetric. We thus proceed to split the sample into three groups, by tercile of magnitude of lagged 24 returns. Figure 6 shows that the dependency of subsequent returns on lagged 24-hour returns is larger for negative, compared with positive returns. This points to an asymmetric effect of momentum. Investors are more likely to extrapolate positive returns than negative returns. On the other hands, there is no pattern for lag 22 return coefficients.

Fourth, we study the heterogeneity of the effect across the different months of the year (Fig. 7), day of the week (Fig. 8), and hour of the day (Fig. 9). While we do find some variation across different times of the day, week, and month, there is no clear pattern emerging. This is consistent with crypto exchanges being integrated around the world, thus differences in time-zones, and holidays seasons smooth the heterogeneity of the effects across time.

Finally, if investors tend to extrapolate past returns, we should expect that trading volume be higher when the absolute value of lagged-24 returns is large. We thus run equation 1), but as outcome variable we use the hourly trading volume. Figure 10 reports the results. It is well known that volume and volatility are correlated, and persistent, and thus it is not surprising to find that lagged absolute returns, a measure of volatility, and subsequent volume are positively correlated. The plot shows that this dependents decays with the lags but rises during the lag 24-hour absolute return. That is, the plot suggests that the response to the lagged 24-hour absolute returns is abnormally high, relative to the general pattern.

V. Trading Strategy

The evidence presented so far shows that crypto returns are serially correlated. We thus proceed to assess the performance of a trading strategy that exploit such predictability of returns. In order to implement a long-short strategy, for each hour we need to have an a sufficient cross section of cryptocurrencies. We thus limit the sample to hours where there are at least 5 cryptocurrencies traded. In the first three years of the sample period, there are fewer than 5 cryptocurrencies traded, so for this part of the paper, the sample period starts in August 2018.

We first examine the average hourly returns based on crypto quintile sorts. More specifically, for each hour, we split the sample into quintiles of lagged 24 returns. We then compute the mean time t return for each quintile. Finally, we compute the arithmetic mean and standard deviation of returns of each quintile across time. Panel A of Table 3 shows the annualized returns, the annualized standard deviation, and the Sharpe ratio. We find a monotonic relationship between the average hourly return and lagged 24-hour return showing that returns are substantially higher following the removal of lower lagged returns (1.84%) compared with higher returns (-0.99%). This pattern is not observed for the placebo lagged 22-hour sorts. A trading strategy that goes long on the first quintile, and short on the fifth quintile would earn an annual return before transaction costs of 258%, with a Sharpe ratio of 2.67.

Based on this, we develop a simple trading strategy that trades hourly based on lagged 24-hour return. Specifically, for each hour, we assign the two cryptocurrencies with the highest lagged returns to the short portfolio and the two with the lowest lagged returns to the long portfolio. The strategy rebalances hourly. Panel B of Table 3 shows that the long-short strategy returns are economically large, at 538% annually. Second, while the strategy is very volatile, its Sharpe ratio is high at 3.52. Third, decomposing the strategy performance we find that most of it comes from the short leg. While the long side produces 144% per year, it is not very different from the average returns on Bitcoin, during our sample period of 2018-2021, of 86% per year. While our sample period is relatively short, there is no evidence that the strategy performance decays over time; if anything, the last two years produce higher performance than the first two years. Finally, a placebo strategy that trades based on lagged 22 hour returns does not yield positive returns or Sharpe ratio. Figure 11 display the cumulative performance of the long-short strategy, as well as the long and short legs relative to the Bitcoin price performance. We also add to the plots the performance of the placebo trading strategy. The long-short strategy shows strong and consistent performance over time (top panel).

While the trading strategy seems highly profitable, it becomes unprofitable once commissions and price impact are taken into account. The strategy requires to turn over the portfolio twice every hour. Fees in crypto exchanges range from 10 to 50bps. Even using the lowest fees, it would cost 20bps per hour, or 4.9% per day. And this is even before considering the price impact of the trading strategy.

VI. Conclusions

This paper makes use of a novel natural experiment afforded by the convention adopted in crpto markets to display daily returns as the returns observed over the previous 24-hours. This convention gives rise to a mechanical and *predictable* shock to perceived returns.

We find robust and economically large evidence consistent with investors not realizing that part of the innovation to these reported daily returns is fully predictable since it removes the returns at the beginning of the sample period as time goes by. That is, investors appear to respond positively to innovation in these returns, which we refer to as “pure momentum”.

While crypto markets may be different from other markets, partially because of the dominance of retail traders, they are large and garner interest on their own right. Our direct evidence documents a novel form of momentum in large, and very liquid, markets. Beyond the crypto market, it suggests a potential explanation for the large periodic patterns of equity returns over salient return intervals. Finally, the always-on nature of crypto markets is viewed by some as a model that should be adopted and expanded to equity markets. For example, 24X National Exchange LLC has applied for a licence to operate a 24-hour stock exchange in the US. Our results imply that the design of salient information in such a market is an important consideration for the welfare of retail traders.

REFERENCES

Asness, Clifford S., Tobias J. Moskowitz, and Lasse Heje Pederson. 2013. "Value and Momentum Everywhere." *The Journal of Finance*, 68(3): 929–985.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer. 2019. "Diagnostic expectations and stock returns." *The Journal of Finance*, 74(6): 2839–2874.

Daniel, Kent, and Tobias J. Moskowitz. 2016. "Momentum crashes." *Journal of Financial Economics*, 122(2): 221–247.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam. 1998. "Investor Psychology and Security Market under- and Overreactions." *The Journal of Finance*, 53(6): 1839–1885.

Grobys, Klaus, and Niranjan Sapkota. 2019. "Cryptocurrencies and momentum." *Economics Letters*, 180(C): 6–10.

Heston, Steven L., and Ronnie Sadka. 2008. "Seasonality in the cross-section of stock returns." *Journal of Financial Economics*, 87(2): 418–445.

Jegadeesh, Narasimhan. 1990. "Evidence of Predictable Behavior of Security Returns." *The Journal of Finance*, 45(3): 881–898.

Jegadeesh, Narasimhan, and Sheridan Titman. 1993. "Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency." *The Journal of Finance*, 48(1): 65–91.

Jia, Boxiang, John W. Goodell, and Dehua Shen. 2022. "Momentum or reversal: Which is the appropriate third factor for cryptocurrencies?" *Finance Research Letters*, 45(C).

Phillips, Blake, Kantura Pukthuanthong, and P. Raghavendra Rau. 2016. "Past Performance May Be an Illusion: Performance, Flows, and Fees in Mutual Funds." *Critical Finance Review*, 5(2): 351–398.

Shen, Dehua, Andrew Urquhart, and Pengfei Wang. 2020. "A three-factor pricing model for cryptocurrencies." *Finance Research Letters*, 34(C).

Sias, Richard. 2007. "Causes and Seasonality of Momentum Profits." *Financial Analysts Journal*, 63(2): 48–54.

VII. Tables and Figures

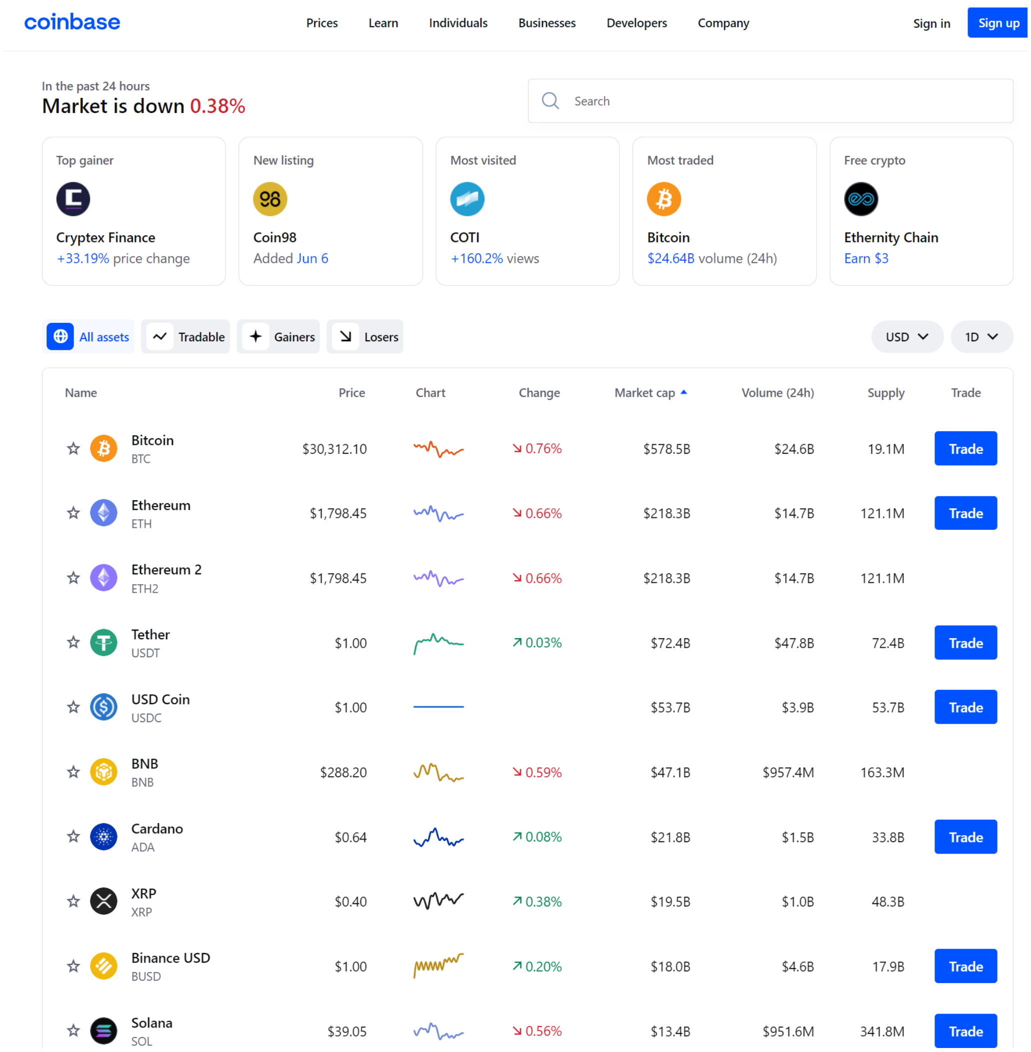

FIGURE 1. DISPLAY OF MARKET RETURNS

Note: This figure is a screenshot of market snapshot taken from Coinbase.com.

Exchange	Reference Price
Binance	Lag 24-Hour
Bitfinex	Lag 24-Hour
Bitstamp	Lag 24-Hour
Coinbase	Lag 24-Hour
crypto.com	Lag 24-Hour
e-toro	12am UTC
FTX	Lag 24-Hour
Gemini	Lag 24-Hour
Kraken	Lag 24-Hour
KuCoin	Lag 24-Hour
Robinhood	12am Local Time
Sofi	12am US ET

Table 1—Trading Venues Reference Price

Note: This table lists the largest crypto trading venues, and the reference price used to compute and present daily returns.

	N. Cryptos	N. Obs.	Hourly Returns Mean	Hourly Returns St. Dev.	Hourly Volume Mean	Hourly Volume St. Dev.
2015	1	3,879	0.014%	0.766%	132,137	480,933
2016	2	12,632	0.008%	0.703%	105,354	128,360
2017	4	23,968	0.052%	1.530%	2,543,533	6,763,297
2018	6	38,909	-0.012%	1.298%	2,057,613	4,200,780
2019	15	89,130	-0.001%	1.045%	679,364	2,677,237
2020	38	189,898	0.014%	1.465%	752,609	2,856,371
2021	138	633,142	0.018%	1.827%	2,065,599	8,517,604
Total	138	991,558	0.015%	1.665%	1,668,238	7,123,576

Table 2—Summary Statistics

Note: This table presents summary statistics of the variables used in the paper.

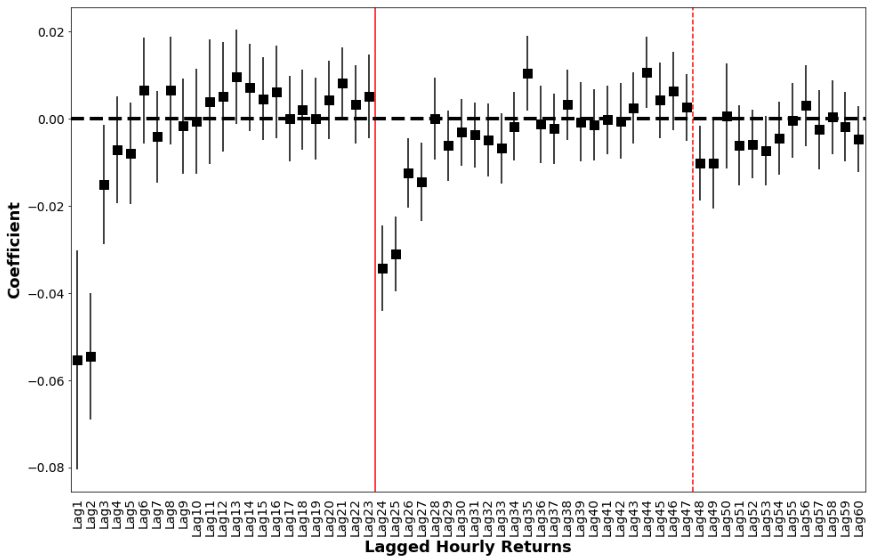

Figure 2. Boxplot of coefficients of hourly lagged returns regression.

Note: The figure presents the lagged return coefficients of equation 1, where the outcome variable is the hourly return of a cryptocurrency, and the variables of interest are the lagged 1-60 hourly returns. The marker represents the estimated coefficient, and the line represents the 95% confidence interval.

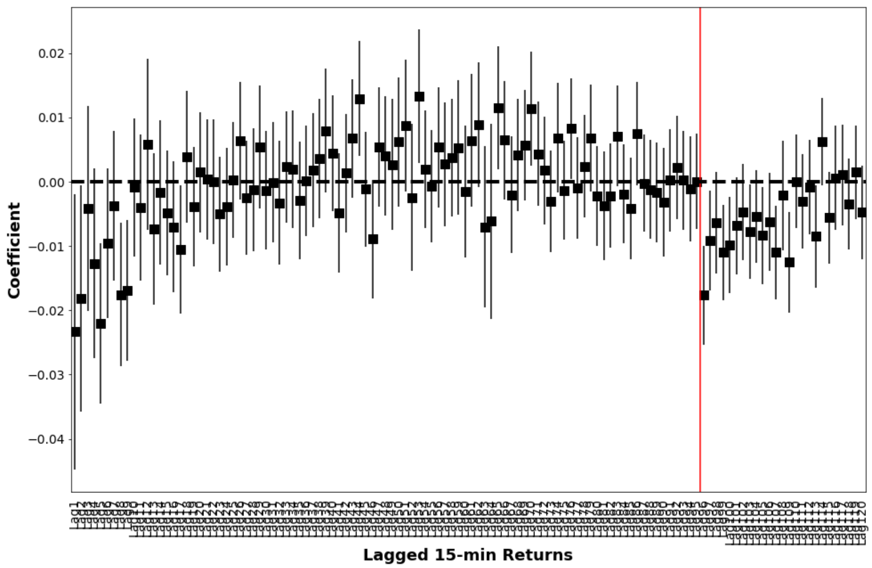

FIGURE 3. BOXPLOT OF COEFFICIENTS OF 15-MINUTE LAGGED RETURNS REGRESSION.

Note: The figure presents the lagged return coefficients of equation 1, where the outcome variable is the 15-minute return of a cryptocurrency, and the variables of interest are the lagged 1-120 15-minutes returns. The marker represents the estimated coefficient, and the line represents the 95% confidence interval.

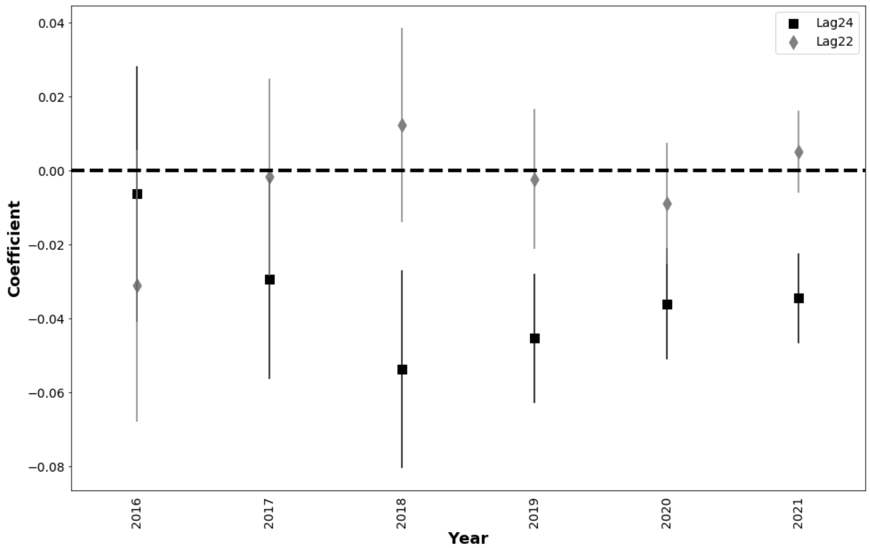

FIGURE 4. BOXPLOT OF 24-HR LAGGED RETURN COEFFICIENT ACROSS YEARS.

Note: The figure presents the lagged 22 and 24 return coefficients of equation 1, where the outcome variable is the hourly return of a cryptocurrency, and the variables of interest are the lagged 1-60 hourly returns for each year in the sample. The marker represents the estimated coefficient, and the line represents the 95% confidence interval.

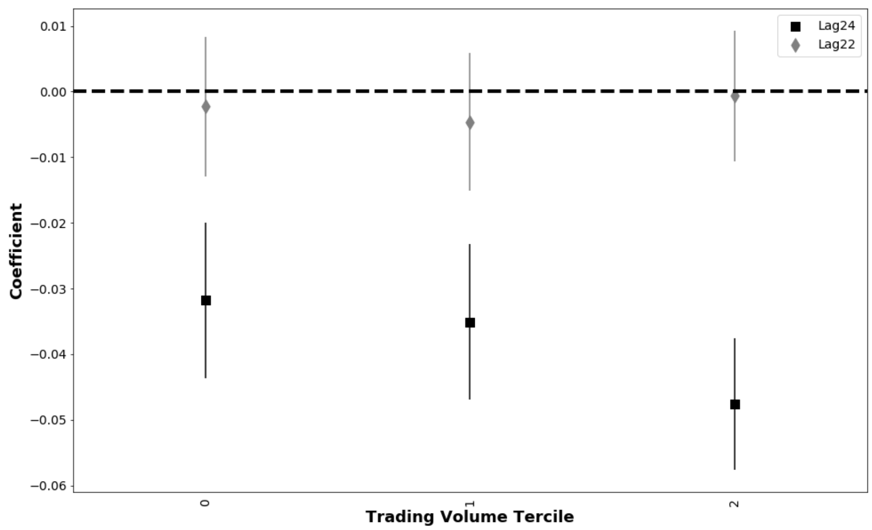

Figure 5. Boxplot of 24-hr and 22-hr lagged return coefficient across trading volume tercile.

Note: The figure presents the lagged 22 and 24 return coefficients of equation 1, where the outcome variable is the hourly return of a cryptocurrency, and the variables of interest are the lagged 1-60 hourly returns, for each tercile of trading volume. The marker represents the estimated coefficient, and the line represents the 95% confidence interval.

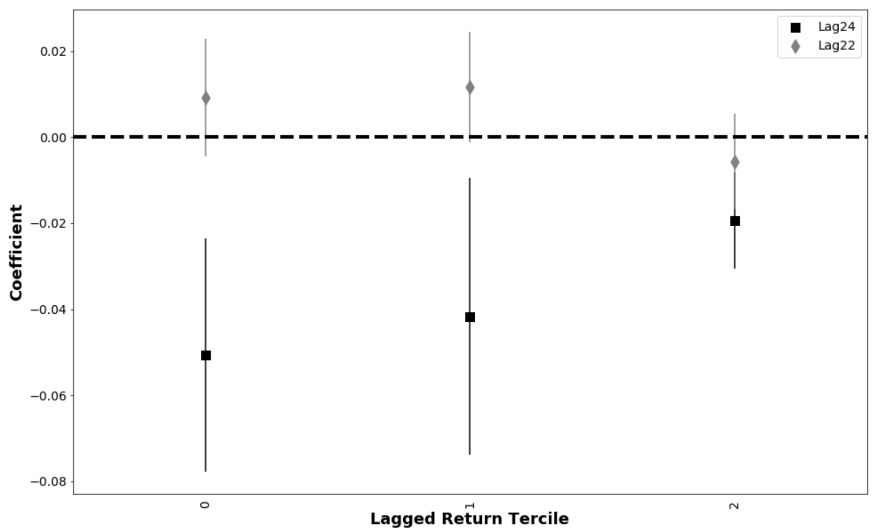

FIGURE 6. BOXPLOT OF 24-HR AND 22-HR LAGGED RETURN COEFFICIENT ACROSS LAGGED RETURN TERCILES.

Note: The figure presents the lagged 22 and 24 return coefficients of equation 1, where the outcome variable is the hourly return of a cryptocurrency, and the variables of interest are the lagged 1-60 hourly returns, for each tercile of lagged 22 and 24 returns. The marker represents the estimated coefficient, and the line represents the 95% confidence interval.

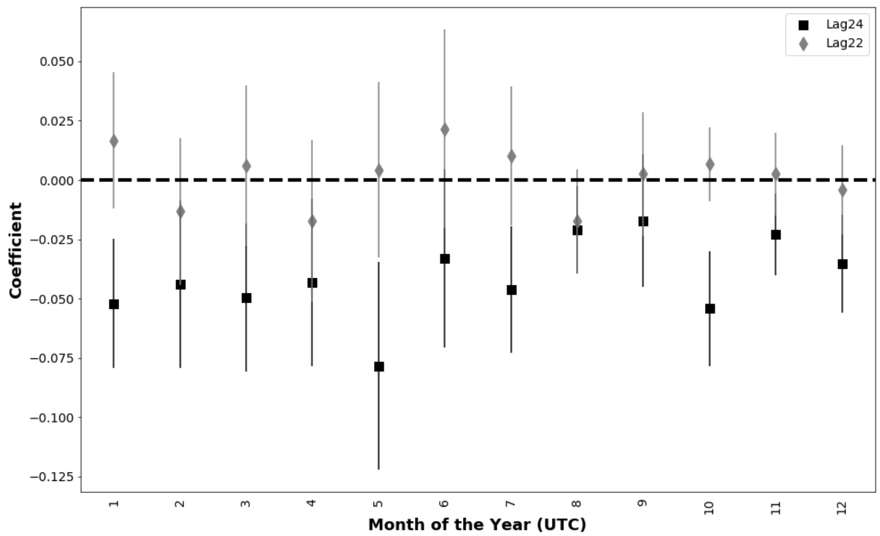

FIGURE 7. BOXPLOT OF 24-HR LAGGED RETURN COEFFICIENT ACROSS MONTHS OF THE YEAR.

Note: The figure presents the lagged 22 and 24 return coefficients of equation 1, where the outcome variable is the hourly return of a cryptocurrency, and the variables of interest are the lagged 1-60 hourly returns, for each month of the year. The marker represents the estimated coefficient, and the line represents the 95% confidence interval.

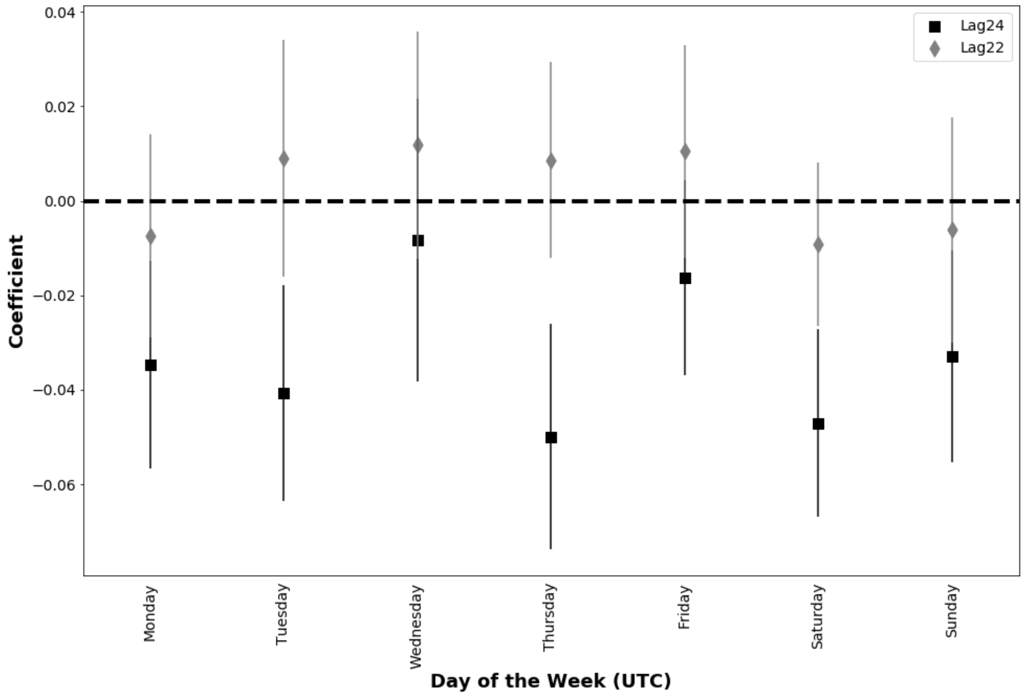

Figure 8. Boxplot of 24-hr lagged return coefficient across days of the week.

Note: The figure presents the lagged 22 and 24 return coefficients of equation 1, where the outcome variable is the hourly return of a cryptocurrency, and the variables of interest are the lagged 1-60 hourly returns, for each day of the week. The marker represents the estimated coefficient, and the line represents the 95% confidence interval.

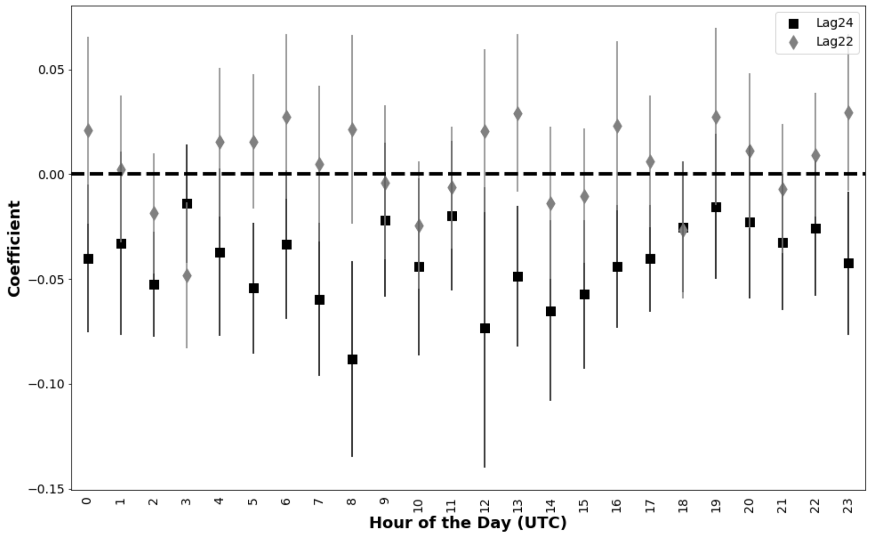

FIGURE 9. BOXPLOT OF 24-HR LAGGED RETURN COEFFICIENT ACROSS HOURS OF THE DAY.

Note: The figure presents the lagged 22 and 24 return coefficients of equation 1, where the outcome variable is the hourly return of a cryptocurrency, and the variables of interest are the lagged 1-60 hourly returns, for each hour of the day. The marker represents the estimated coefficient, and the line represents the 95% confidence interval.

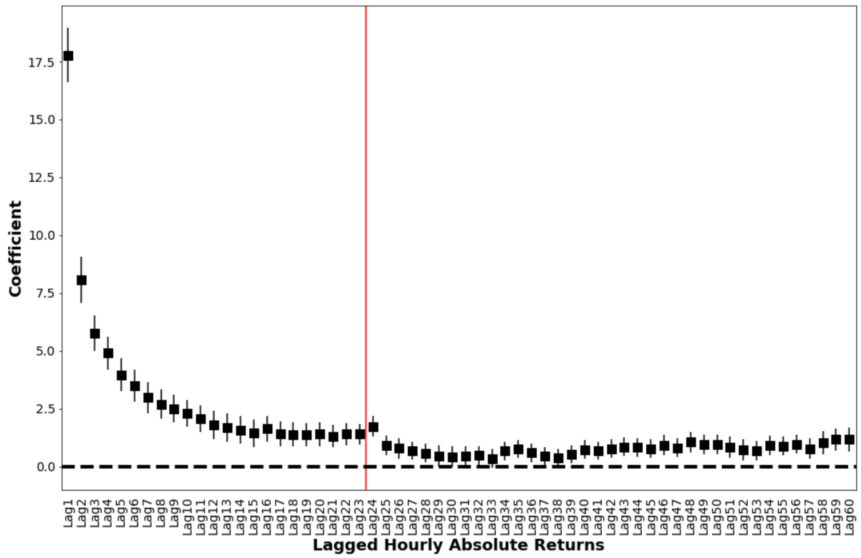

Figure 10. Trading volume and lagged absolute returns.

Note: The figure presents the lagged return coefficients of equation 1, where the outcome variable is the hourly trading volume of a cryptocurrency, and the variables of interest are the lagged 1-60 hourly absolute returns. The marker represents the estimated coefficient, and the line represents the 95% confidence interval.

Panel A: Quintile Returns

	Quintile	Ret	Stdev	Sharpe
Lag24	1	1.84	1.05	1.75
	2	1.59	1.0	1.58
	3	1.17	1.01	1.16
	4	0.97	1.0	0.97
	5	-0.99	1.08	-0.92
	Q1-Q5	2.58		2.67
Lag22	1	0.11	1.06	0.11
	2	1.53	1.02	1.5
	3	1.27	0.99	1.28
	4	0.98	1.0	0.98
	5	0.82	1.08	0.76
	Q1-Q5	0.71		-0.65

Panel B: Long-Short Strategy

	Year	Long-Short Strategy			Long Strategy			Short Strategy			BTC Buy-and-Hold		
		Ret	Stdev	Sharpe	Ret	Stdev	Sharpe	Ret	Stdev	Sharpe	Ret	Stdev	Sharpe
Lag24	2018	0.13	0.72	0.19	-2.14	1.08	-1.97	2.27	1.06	2.15	-1.02	0.73	-1.4
	2019	2.63	0.87	3.04	1.42	0.91	1.56	1.22	0.89	1.37	0.77	0.69	1.12
	2020	6.42	1.49	4.3	3.22	1.34	2.41	3.21	1.43	2.25	1.7	0.77	2.2
	2021	9.13	2.17	4.21	1.1	1.83	0.6	8.03	1.92	4.19	0.84	0.87	0.97
	Total	5.38	1.53	3.52	1.44	1.38	1.04	3.95	1.43	2.76	0.86	0.78	1.11
Lag22	2018	-0.35	0.73	-0.48	-2.01	1.08	-1.85	1.66	1.07	1.55	-1.02	0.73	-1.4
	2019	-0.58	0.84	-0.69	-0.59	0.86	-0.69	0.01	0.89	0.01	0.77	0.69	1.12
	2020	-1.61	1.57	-1.03	-0.22	1.39	-0.16	-1.39	1.42	-0.98	1.7	0.77	2.2
	2021	-2.6	2.2	-1.18	-1.54	1.84	-0.84	-1.06	1.9	-0.56	0.84	0.87	0.97
	Total	-1.46	1.56	-0.93	-0.93	1.39	-0.67	-0.53	1.42	-0.37	0.86	0.78	1.11

TABLE 3—TRADING STRATEGY.

Note: This table presents the results of two trading strategies. We only keep data where hourly returns are present for at least 5 cryptos. Each hour, we rank cryptos according to their lagged 24 or 22 hourly returns. In Panel A, each hour we classify crypto into quintiles according to their ranks. We then compute the average annualized return, standard deviation, and Sharpe ratio for each quintile. In Panel B, each hour we go long the two cryptos with the lowest returns, and go short the two cryptos with the highest returns.

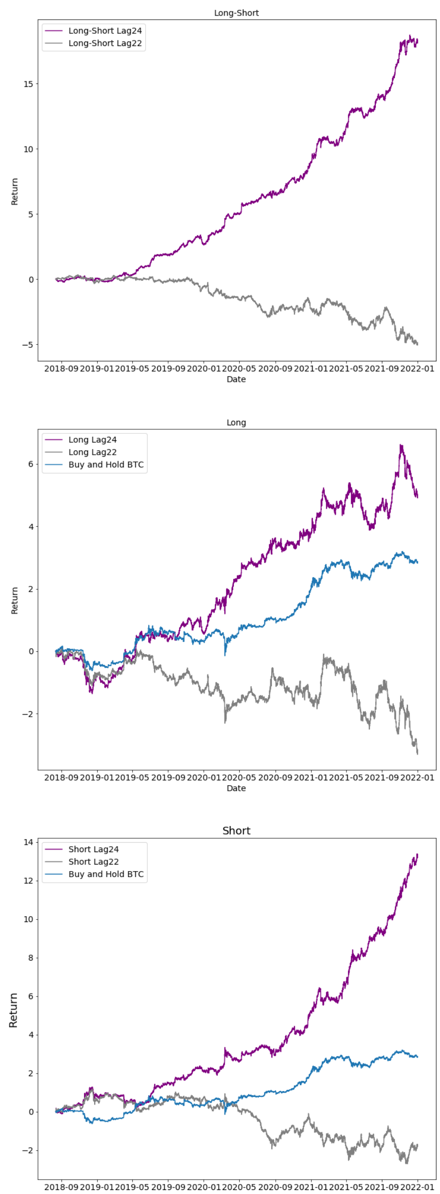

FIGURE 11. LONG-SHORT TRADING STRATEGY

Note: The figure presents the cumulative returns of a trading strategy each hour we go long the two cryptos with the lowest lagged 22 or 24 returns, and go short the two cryptos with the highest lag 22 or 24 returns. The top/middle/bottom sub-figures shows the cumulative return of the long-short/long-only/short-only strategies.