Machine Learning and the Cross-Section of Cryptocurrency Returns

Nusret Cakici†, *Syed Jawad Hussain Shahzad*‡, *Barbara Będowska-Sójka*§, *Adam Zaremba*δ⁑

Abstract

We employ a repertoire of machine learning models to explore the cross-sectional return predictability in cryptocurrency markets. While all methods generate substantial economic gains, those that account for nonlinearities and interactions fare the best. The return predictability derives mainly from a handful of simple features—such as idiosyncratic volatility, past alpha, or maximum daily return—and is likely driven by mispricing. Accordingly, abnormal returns originate predominantly from short positions, concentrate in hard-to-arbitrage assets, and gradually decline over time. Despite a high portfolio turnover, machine learning strategies remain a profitable net of trading costs. However, they critically depend on shorting small cryptocurrencies, which may pose challenges in practice.

Keywords: cryptocurrency markets, machine learning, return predictability, limits to arbitrage, asset pricing, the cross-section of returns

JEL codes: G11, G12, G17

This version: 6 December 2022

* Barbara Będowska-Sójka acknowledges the support of the National Science Center of Poland [Grant No. 2021/41/B/HS4/02443].

⁑ Corresponding author.

† Nusret Cakici, Gabelli School of Business, Fordham University, 45 Columbus Avenue, Room 510, New York, NY 10023, USA, tel. +1-212-636-6776, email: cakici@fordham.edu.

‡ Syed Jawad Hussain Shahzad, Chair of Social & Sustainable Finance, Montpellier Business School, 2300 Avenue des Moulins, 34080 Montpellier, France, j.syed@montpellier-bs.com.

§ Barbara Będowska-Sójka, Department of Econometrics, Institute of Informatics and Quantitative Economics, Poznan University of Economics and Business, al. Niepodległości 10, 61-875 Poznań, Poland, Barbara-bedowska-sojka@ue.poznan.pl

δ Adam Zaremba, 1) Montpellier Business School, 2300 avenue des Moulins, 34185 Montpellier, France, a.zaremba@montpellier-bs.com; 2) Department of Investment and Capital Markets, Institute of Finance, Poznan University of Economics and Business, al. Niepodległości 10, 61-875 Poznań, Poland, adam.zaremba@ue.poznan.pl.

1. Introduction

Cryptocurrency literature has documented a growing list of characteristics that predict cross-sectional returns. Some of them—such as momentum, size, or reversal (Liu et al., 2022; Bianchi et al., 2022)—are parallels of similar phenomena in equity markets; others, such as network activity measures (Liu & Tsyvinski, 2021; Cong et al., 2022), are inherently specific to cryptocurrencies. The emerging cryptocurrency "factor zoo"—comprising potentially noisy and correlated predictors—may require methods beyond simple portfolio sorts or cross-sectional regressions. Recent advances in machine learning methods appear to be a natural response to this task. Due to their capacity to handle vast multidimensional data, select best predictors, and account for nonlinearities and interactions (Gu et al., 2020; Giglio et al., 2022), machine learning models are well poised to face the cryptocurrency landscape.

In this paper, we combine machine learning with asset pricing research in order to gain new insights into cross-sectional return predictability in cryptocurrency markets. Building on a comprehensive dataset that covers more than 500 major coins and tokens listed across 250 exchanges over the years of 2017 to 2022, we identify, classify, and reproduce 34 cryptocurrency characteristics. Next, we use these to feed 10 popular machine learning models; this includes dimension reduction techniques, raw and regularized regressions, tree methods, and neural networks. Finally, we explore their prediction performance. The aim of these exercises is not only to assess the machine learning effectiveness in the cryptocurrency world but—first and foremost—to enrich our understanding of pricing mechanisms within this new and growing asset class.

Our findings contribute in four ways. First, we demonstrate that machine learning techniques can be successfully deployed to predict the cross-section of cryptocurrency returns. While all our forecasting models generate substantial economic gains, the tree methods and neural networks prove particularly effective. Their superiority derives from the ability to capture interactions and nonlinearities in returns; this highlights the complexity of asset pricing in cryptocurrency markets. The linear models, which cannot capture these phenomena, typically lag.

Besides the mere prediction accuracy, the machine learning models can be forged into successful investment strategies via portfolio sorts. The best-performing method is the forecast combination, which equally weights outputs from individual models. In a nutshell, while all models have their benefits, merging them works best. A long-short value-weighted combination strategy yields a mean weekly return of 4.72% at an annualized Sharpe ratio of 5.37. The abnormal returns survive controlling for common factor exposures and hold for alternative weighting schemes. Furthermore, as seen in Coqueret (2022), investors typically benefit from using longer-horizon forecasts. As a result, strategies based on daily forecast yield lower—though still significant—abnormal returns than either weekly or bi-weekly portfolios.

Importantly, because space constraints prevent us from displaying the full details of our machine leaning strategies, we supplement the paper with a visual tool to view them on a dedicated website.[1] It allows researchers to compare the performance of all strategies discussed in this study.

Second, our findings highlight the main drivers of the cross-sectional variation in cryptocurrency returns. The machine learning models allow for pinpointing the essential return predictors among the "factor zoo." Our variable importance analysis suggests that cryptocurrency returns are determined by a handful of uncomplicated signals—including idiosyncratic volatility, CAPM alpha, maximum daily return, nominal price, value at risk, and distance to a 90-day high. Notably, all these variables originate from simple data types—such as prices and returns. Moreover, when predictors are grouped into general economic categories, what matters the most are past returns, volatility, and liquidity.

Our third contribution concerns the sources of cross-sectional return predictability in cryptocurrency markets. A deeper look at machine learning forecasts provides insights into the nature of asset pricing mechanisms. Our findings align most closely with the mispricing explanation of return predictability: market inefficiencies emerge due to, e.g., investors' limited rationality and persistence if they cannot be easily arbitraged away. Consistent with this, machine learning alphas predominantly come from the short legs of long-short strategies. A similar pattern also characterizes many equity anomalies (Stambaugh et al., 2012, 2015; Avramov et al., 2013, 2019), as short selling costs and constraints may impede attempts to eliminate mispricing.

Moreover, the abnormal returns on machine learning portfolios mainly concentrate in difficult-to-trade cryptocurrencies. While our strategies are profitable everywhere, they thrive particularly in market segments that are characterized by excessive limits to arbitrate: high idiosyncratic risk, illiquidity, and bid-ask spread. In other words, the return predictability exists to the extent that impediments to arbitrage prevent traders from fully eliminating inefficiencies. Overall, the alphas in the tercile of assets with the highest limits to arbitrate are roughly three times as high as in the tercile of easiest-to-trade cryptocurrencies.

Last, in addition to the cross-sectional variation above, the magnitude of return predictability also varies over time. In particular, the abnormal returns on machine learning strategies gradually decline over time—as the markets become more efficient. Whereas their alphas have remained significant even in the second half of our sample (April 2020 to July 2022), their sheer size has declined by about 30%.

[1] Https://azaremba.shinyapps.io/Cryptocurrency_machine_learning/.

Our fourth—and final—contribution concerns an investor's practical perspective. The cryptocurrency machine learning strategies produce phenomenal alphas, which can hardly be encountered in other asset classes. In particular, their Sharpe ratios visibly beat those of machine learning strategies in developed or emerging stock markets (e.g., Gu et al., 2020; Liu et al., 2022). Can these profits be harvested in practice? We find that despite a sizeable portfolio turnover, machine learning strategies survive net of transaction costs—even under restrictive assumptions. Admittedly, the trading costs can consume as much as 60% of the payoffs of long-short machine learning strategies. Nonetheless, the remaining profits are still sizeable and are at par with their counterparts in equity markets.

However, the cryptocurrency machine learning strategies come with two caveats. First, they exhibit substantial tail risk; therefore, an investor must be prepared for significant drawdowns of up to 30-50% of the invested capital. Second, they require taking massive short positions in small capitalizations cryptocurrencies. In particular, the last point may be challenging or even unfeasible. A successful implementation of machine learning in cryptocurrency markets depends on sorting out these two issues.

Our study connects with three main strains of finance research. First, we extend the discussion on cross-sectional return predictability in cryptocurrency markets. A growing list of studies has demonstrated dozens of signals that capture the cross-sectional variation in cryptocurrency returns. Examples include—but are not limited to—momentum, size, liquidity, reversal, downside risk, crypto-specific network and on-chain measures, as well as macroeconomic exposures (e.g., Bhambhwani et al., 2019; Liu et al., 2021, 2022; Liu & Tsyvinski, 2021; Zhang et al., 2021; Borri et al., 2022; Cong et al., 2022; Bianchi et al., 2022). Against this backdrop, our paper is most closely linked to studies that attempt to aggregate the multidimensionality of returns into a uniform pricing model. For example, Liu et al. (2022) propose a three-factor model, and Babiak and Bianchi (2022) implement the instrumented principal component analysis of Kelly et al. (2019).

Second, our paper adds to the research on machine learning applications to cross-sectional return predictability within financial markets. Earlier studies concentrated on diverse asset classes, including U.S. equities (Freyberger et al., 2020; Gu et al., 2020, 2021; Avramov et al., 2021; Coqueret, 2022), international stocks (Leippold et al., 2021; Drobetz & Otto, 2021; Tobek & Hronec, 2021; Azevedo et al., 2022; Cakici et al., 2022; Fieberg et al., 2022; Hanauer & Kalsback, 2022), corporate bonds (Bali et al., 2022), U.S. Treasury bonds (Bianchi et al., 2021), and commodities (Struck & Cheng, 2020; Rad et al., 2021). Cryptocurrencies have thus remained a largely uncharted territory.

Third, our findings relate to the evidence on economic constraints on return predictability. Existing literature concentrates mainly on equities. For example, Hong et al. (2000), Fama and French (2008, 2012), and Hou et al. (2020) show that anomalies concentrate primarily in small and illiquid securities. Pontiff (2006) and McLean (2010) accentuate the role of idiosyncratic volatility limiting arbitrage. Stambaugh et al. (2012, 2015) and Avramov et al. (2013, 2019) argue that most abnormal returns come from short legs of anomaly portfolios. Novy-Marx and Velikov (2016, 2019), Patton and Weller (2020), and Chen and Velikov (2021) document the sensitivity of profits from cross-sectional return predictions to trading costs. Finally, abundant literature—encompassing the works of Schwert (2007), Chordia et al. (2014), McLean and Pontiff (2016), Caluzzo et al. (2019), and Linnainmaa and Roberts (2018)—demonstrate the gradual decrease in return predictability over time.[2] Against this background, our article is particularly associated with Avramov et al. (2022)—who explore economic restrictions on machine learning implementation effectiveness in equity markets. Similarly, they conclude that machine learning methods extract profitability mainly from difficult-to-arbitrage securities.

The remainder of the study proceeds as follows. Section 2 discusses our data and findings. Section 3 presents the baseline empirical findings on the prediction performance of machine learning models, as well as the contribution of individual crypto characteristics. Section 4 concentrates on machine learning portfolios. Section 5 scrutinizes economic restrictions on machine learning strategies. Section 6 concerns the practical investor perspective. Finally, Section 7 concludes the study.

2. Data and Methods

This section summarizes the data and methods used in this study. We begin by presenting our dataset and the sample of cryptocurrency characteristics. We then describe the machine learning models employed, as well as the methods of evaluating their performance.

2.1. Data Sources and Sample Preparation

Following Bianchi and Babiak (2022), we combine data from several databases to maximize market coverage and availability of cryptocurrency characteristics. Specifically, we source OHLC price and volume data from CryptoCompare.com and blockchain activity from IntoTheBlock.com. Unlike many other data providers, CryptoCompare aggregates transaction data from more than 250 centralized exchanges around the world to provide an accurate price estimation. The market data is volume-weighted across the exchanges, linking the data source's importance with trading activity. As a result, smaller trading venues are given less emphasis in the aggregation process than the big and liquid markets. CryptoCompare has been recommended by Alexander and Dakos (2019) due to its superior reliability; it is frequently employed as a prime data source (Borri, 2019; Lucchini et al., 2020; Bianchi et al., 2022; Borri & Shankhnov, 2022). On the other hand,

[2] Notably, several studies—including Jacobs (2016) and Jacobs and Müller (2020)—contest the conclusions on declining return predictability.

IntoTheBlock—used by Cong et al. (2022), Bianchi and Babiak (2022), and Hoang and Baur (2022), among others—is one of the most common sources of on-chain activity.

Cryptocurrency tickers used by various providers may differ. Hence, in order to ensure correct matching coins from CryptoCompare and IntoTheBlock, we use a two-step procedure. In the first pass, we pair the cryptocurrencies by their full name rather than by ticker. However, cryptocurrency names may also exhibit minor discrepancies across databases. Hence, in the second pass, we turn to matching by tickers. We only retain these assets with market data available from both sources and require price consistency. Precisely, we discard coins with identical tickers but prices differing on average by more than 5%. In addition, we verify that their names are qualitatively consistent on a case-by-case basis.

We clean our sample with a series of filters to eliminate potential data errors, mainly following Bianchi and Babiak (2022). We begin with static screens. We exclude (i) stablecoins, e.g., USDT or DAI (either centralized or algorithmically stabilized); (ii) coins backed by or tracking prices of precious metals, such as gold; (iii) cryptocurrencies serving as collaterals for derivative platforms, e.g., SNX; and (iv) so-called "wrapped coins," including Wrapped Bitcoin (WBTC). The cryptocurrency classification for these filters is based on CoinMarketCap.com.

Next, we also apply a battery of dynamic filters. First, following Cong et al. (2022), we delete any observations with non-positive trading volume, market capitalization, or price. Second, as in Bianchi et al. (2022), we mitigate issues with "fake" or "erroneous" volume by dropping observations with the ratio of daily traded volume to market capitalization exceeding one. Third, to align our study with practice, we exclude the smallest cryptocurrencies—which may be difficult to trade. To be precise, we adopt the approach of Liu et al. (2022): setting a minimum threshold of $1 million for cryptocurrency market capitalization. Finally, as in Bianchi and Babiak (2022), we discard extreme returns as they may likely originate from database errors. Concretely, we remove daily returns below -99.9% and above 100%—as well as weekly returns below -99.9% and exceeding 200%.

After we apply all the filters, our sample contains 574 distinct cryptocurrencies. The study period runs from July 1, 2017 to July 6, 2022 and the start time is defined as 00:00:00 UTC. Obviously, the number of coins is not stable over time and grows gradually along with the evolution of the cryptocurrency markets. Figure 1 provides a snapshot of the research sample. As seen in Panel A, the number of available assets increased from 30 in July 2017 to reach 300 in late 2018. Notably, the cross-section may seem modest when compared to the existing cryptocurrency universe; at the time of writing, CoinMarketCap.com informs more than 24 thousand coins traded at different exchanges. Nevertheless, by focusing on larger and more tradeable tokens, we simultaneously capture most of the global market capitalization. As seen in Panel C, the total market value of the sample is below less than $100 billion in 2017 to exceed two trillion in 2021—effectively capturing most of the tradeable cryptocurrency universe.

[Insert Figure 1 about here]

Our study period—although relatively short—covers a diverse set of market, economic, and geopolitical circumstances; this includes the COVID-19 pandemic outbreak, the subsequent global recession, the stock market downturn and rally in 2020, and the Russian invasion of Ukraine. It also captures various institutional and regulatory changes, such as introducing novel instruments—Chicago Mercantile Exchange Bitcoin and Ethereum futures or Bitcoin ETFs—or the crypto exchange ban in China. Consequently, it reflects a broad array of conditions in the cryptocurrency market—including the initial coin offering frenzy in 2017, the subsequent "crypto-winter," as well as massive cryptocurrency crash at the onset of the pandemic in 2020. In summary, our research sample comprehensively represents conditions and tendencies in cryptocurrency markets.

Finally, Panel B of Figure 1 illustrates the cross-sectional composition of our research sample. An average token in our selection has 774 trading days' worth of data.[3] A total of 89 coins have a history longer than 1,500 trading days. On the other hand, the sample also includes a substantial representation of cryptocurrencies' relatively short history, and 186 coins have less than one year of data available.

2.2. Cryptocurrency Characteristics

Our machine-learning forecasting models require inputs in the form of coin characteristics. Hence, we identify, classify, and reproduce 34 return-predicting variables from finance literature. We build our selection mainly on Liu et al. (2022) and Bianchi and Babiak (2022), which we complement with more recent findings from cryptocurrency research. We also use these source papers to follow possibly close the calculation procedure. Below, we briefly summarize the predictors, while Table 1 contains a detailed description.

[Insert Table 1 about here]

We group the signals into six broad categories based on their underlying economic intuition. The first group—the measures of *on-chain activity*—encompasses new addresses (*new_add*), active addresses (*active_add*), and network-to-market ratio (*bm*). These signals are sometimes treated as proxies for cryptocurrency value (Pagnotta & Buraschi, 2018; Liu et al., 2021; Cong et al., 2022; Liebi, 2022).

Next, we consider an array of variables associated with *liquidity* that has been explored in cryptocurrency literature (Brauneis et al., 2021; Li et al., 2021; Zhang & Li, 2021; Han et al., 2022; Dong et al., 2022; Liu et al., 2022): trading volume (*volume*), market value

[3] Unlike stocks, cryptocurrencies are traded 365 days a year. As a result, 774 trading days imply roughly 25 months.

(*size*), bid-ask spread (*bidask*), illiquidity ratio (*illiq*), turnover (*turn*), and detrended turnover (*dto*). As in Leivrik (2021), Bianchi and Babiak (2022), and Dong et al. (2022), we also account for measures of variability in liquidity—namely, turnover and trading volume volatility (*std_dto*, *std_vol*), as well as 30- and 60-day volume shocks (*volsh_30d*, *volsh_60d*).

The third category includes various proxies for security-level price volatility and risk that have been found to either positively or negatively correlate with future cryptocurrency payoffs (Jia et al., 2021; Zhang & Li, 2020; Burggraf & Rudolf, 2021; Zhang et al., 2021; Dobrynskaya & Dubrowskiy, 2022): realized volatility (*rvol*), the capital asset pricing model (CAPM), beta (*beta*), idiosyncratic risk (*ivol*), and value-at-risk (*var*).

The fourth group, associated with *past returns*, is relatively broad. It encompasses assorted variables that are derived from past price changes. These signals capture the positive or negative correlation between past and future coin performance (Grobys & Sapkota, 2019; Liu et al., 2020; Shen et al., 2020; Tzouvanas et al., 2020; Dobrynskaya, 2021; Liu & Tsyvinski, 2021; Bianchi et al., 2022; Jia et al., 2022). Specifically, we calculate several momentum measures motivated by Liu et al. (2022): seven-day, 13-day, 23-day, 31-day, and intermediate momentum (*r7_2*, *r13_2*, *r22_2*, *r31_2*, *r30_14*); two reversal signals: daily reversal (*r2_1*) and long-term reversal (*r180_r60*); as well as two other measures: the 90-day high (*90dh*) and the CAPM alpha (*alpha*).

Besides the average past returns, we also study their *distribution*. In addition to skewness (*skew*) and kurtosis (*kurt*), this category also includes the maximum and minimum daily returns (*max*, *min*)—examined initially by Bali et al. (2011). Numerous works (e.g., Grobys & Junttila, 2021; Jia et al., 2021; Lin et al., 2021; Liu et al., 2021; Ozdamar et al., 2021) suggest that such distributional properties may contain information about future coin returns.

Finally, our set also includes several variables classified as *other*—which do not fit well into the categories above. These include the nominal price (*prc*), analyzed by Miller and Scholes (1982); the cross-sectional seasonality (*seas*) from Keloharju et al. (2016); and the variables associated with distribution perception distortions: the salience theory measure (*st*) by Cosemans and Frehen (2021) and the chronological return ordering (*cro*) variables of Mohrschladt (2021). The cryptocurrency counterparts of these anomalies have been scrutinized by Zaremba et al. (2020), Cai and Zhao (2021), Chen et al. (2022), and Liu et al. (2022)—among others.

Table 2 presents the summary statistics for the main variables in the sample. We report the averages, standard deviations, skewness coefficients, and key percentiles. Notably, most variables have asymmetric distributions. Many of them—such as on-chain measures, size, or volume—exhibit positive skewness. This phenomenon reflects the specific market structure, with a small number of very large cryptocurrencies (such as Bitcoin or

Ethereum). On the other hand, multiple liquidity variables display negative skewness—suggesting a long tail of illiquid coins.

[Insert Table 2 about here]

2.3. Machine Learning Models

We build on Gu et al. (2020) and use a general additive prediction model to capture the association between the cryptocurrency returns and their characteristics:

$$r_{i,t+1} = E_t(r_{i,t+1}) + \varepsilon_{i,t+1}, \tag{1}$$

with $r_{i,t+1}$ denoting the excess return on token $i = 1, \ldots, N_T$ in the week (seven days) $t = 1, \ldots, T$. Most cryptocurrency asset pricing studies typically focus on explaining daily (e.g., Bianchi & Babiak, 2022; Bianchi et al., 2022) or weekly (e.g., Cong & He, 2019; Liu & Tsyvinski, 2021; Liu et al., 2022) returns. We opt for the latter, a less granular, approach for two reasons. First, the weekly returns are less noisy—allowing us to focus more on the underlying return drivers. Second, the less frequent portfolio reconstruction is more aligned with an investor's perspective as it generates lower trading costs.

We compute the expected returns $E_t(r_{i,t+1})$ as a constant function of cryptocurrency characteristics:

$$E_t(r_{i,t+1}) = g(z_{i,t}), \tag{2}$$

where $z_{i,t}$ denotes the vector of stock characteristics, which contains the 34 variables described in Section 2.1. The function $g(z_{i,t})$ estimates the rates of return that solely rely on information from token *i* and day *t*. In other words, it does not take into account the characteristics of other cryptocurrencies or from earlier periods. Moreover, its precise form is not specified. The approximated functions differ across the machine learning algorithms and may be parametric or nonparametric, linear or nonlinear.

As in Babiak and Bianchi (2022), we consistently predict returns and construct strategies that assume a skipping in portfolio formation equals one day. In other words, we forecast returns based on data from day *d*-1, use them to form portfolios in *d*, and track their returns over the next week (i.e., from day *d*+1 *to d*+7). This one-day lag not only aligns well with the practical investor perspective, but it also allows for the mitigation of potential microstructure issues that are associated with bid-ask bounce (Zaremba et al., 2021; Bouri et al., 2022).

We train all models to forecast the true returns by minimizing the out-of-sample mean squared forecast error:

$$MSFE_{t+1} = \frac{1}{N_{t+1}} \sum_{i=1}^{N_{t+1}} \left(\hat{\varepsilon}_{i,t+1}\right)^2, \tag{3}$$

with $\hat{\varepsilon}_{i,t+1}$ indicating the individual prediction error for the cryptocurrency i, and N_{t+1} representing the number of tokens at t+1. As a general rule, we aim to find a forecasting model with the best prediction accuracy among a pool of candidates.

Machine learning literature offers a broad—and growing—number of prediction models. Our selection of a sample of representative methods builds on several seminal studies from asset pricing literature; this includes Gu et al. (2020), Bali et al. (2022), and Leippold et al. (2022). In consequence, we consider 10 distinct models—which are briefly summarized below. Regarding further implementation details, we closely reproduce procedures from Cakici et al. (2022).

Ordinary Least Squares Regression (OLS). The OLS regression assumes fitting a multiple predictive regression that utilizes all features as model inputs. The model is relatively simple and does not involve regularizing, hyperparameters tuning, or validation. Nonetheless, it is highly prone to overfitting—especially in a high-dimensional setting (Gu et al., 2020).

Partial Least Squares (PLS). PLS is an effective dimension-reduction technique that accounts for the association between covariates and security returns. Advocated by Kelly and Pruitt (2013, 2015), the three-pass PLS regression concentrates on the features most strongly correlated with cryptocurrency payoffs. Specifically, the individual predictors are aggregated into composite factors to maximize the correlation between them and future returns. Subsequently, these newly created factors are used as inputs to predictive regressions. The tuning parameter is the number of components in the regression.

Penalized Linear Regressions (LASSO, ENET). Penalize linear regressions to cope with the overfitting problem by imposing a penalty term on slope coefficients. We employ two popular regularization approaches: the least absolute shrinkage and selection operator (LASSO), and the elastic net (ENET) of Zou and Hastie (2005). LASSO penalizes the model proportionally to the absolute values of its coefficients. On the other hand, ENET's penalty function combines the component of LASSO and the ridge regression of Hoerl and Kennard (1970)—which concentrates on squared coefficients. In our case, both components are weighted equally. The primary benefit of ENET over LASSO is its effectiveness in dealing with a correlation between covariates (for details, see Zou & Hastie, 2005; Diebold & Shin, 2019). The tuning parameter in both models is the penalty term.

Support Vector Machine (SVM). The support vector machine (SVM) models search for hyperplanes to divide the multidimensional vector space territorially into clusters. In our case, the vectors represent cryptocurrency features. The SVM procedure aims to minimize the number of misclassified vectors, maximizing the distance between the correctly classified ones. SVM may be used for binary and multi-class problems in classification and regression contexts. We estimate SVM parameters using the average stochastic gradient descent method (Xu, 2011).

Tree Models (RF, GBRT). The tree models are flexible non-parametric algorithms that effectively capture returns' interactions and nonlinearities. They divide observations (cryptocurrency characteristics) into distinct subcategories, typically called “leaves.” A tree is constructed in stages, and the splitting variables and decision nodes determine the tree's structure. At each splitting point, a splitting variable generates two disjoint branches. The tree continues to grow until the terminal leaves are reached.

Individual trees tend to overfit the data, so they require heavy regularization. We choose two popular approaches: random forests (RF) and gradient-boosted regression trees (GBRT). RF relies on the bootstrap aggregation algorithm called “bagging” (Breiman, 2001). Specifically, it assumes averaging multiple trees based on bootstrapped subsamples on the original data. The GBRT technique, in turn, assumes fitting subsequent trees based on earlier trees’ residuals and forming ensemble predictions by aggregating numerous trees consecutively multiplied by the learning rate (0.1). We fit the GBRT mode using the least-squares boosting method (Breiman, 2001; Hastie et al., 2008) with between 100 and 200 learning cycles.

Feed-Forward Neural Networks (NN1, NN2). The feed-forward neural networks comprise: (i) an “input layer,” i.e., cryptocurrency characteristics; (ii) a number of “hidden layers” containing activation and functions that transform the features; and (iii) an “output layer” that converts the hidden layers’ outcomes into return predictions. The information flows through the neurons, from the input through hidden to output layer, eventually aggregates into forecasts.

The neural networks allow for the accounting of non-linearities and interactions, and their flexibility increases with the number of layers. We use two different versions: one (NN1) and two (NN2) layers, with the first and second layer comprising 16 and 8 neurons—respectively. Following Gu et al. (2020), we employ the rectified linear unit as the activation function. We optimize the trees using the Adam algorithm of Kingma and Ba (2014).

Forecast Combination. To create an ensemble forecast, the forecast combination (COMB) method weighs a number of predictions from various models. Combining several predictions into one has strong statistical roots (Bates & Granger, 1969; Clemen, 1989; Timmermann, 2006). Merging the models helps diminish their variance, eventually reducing the forecast error (Petropoulos et al., 2022). As in Bali et al. (2022), we obtain the COMB forecast by calculating a simple equal-weighted average of the predictions from nine individual models discussed above: OLS, PLS, LASSO, ENET, SVM, RF, GBRT, NN1, and NN2.

We estimate the model's parameters, tune their hyperparameters, and assess their performance using standard methods from machine learning literature. We split the sample into three consecutive subperiods, keeping their temporal order: a training period of 150 days, a validation period of 50 days, and the testing period comprising the subsequent week. Hence, the first testing period runs from July 1, 2017 to November 27, 2017. The training sample is used to estimate model parameters subject to pre-specified hyperparameters (specific to a given model type). Next, we use the subsequent validation period to optimize the hyperparameters to minimize the objective loss function. Last, we test the model's accuracy using the following week. The testing week is not included in the training or validation samples. Finally, we re-estimate the model each week and reproduce this procedure until we reach the end of the entire study period. We assume a fixed (or rolling) training window, which means that the length of the training, validation, and testing samples is held constant and rolled forward at each reestimation.

2.4. Performance Evaluation

Our baseline measure of forecasting accuracy is the out-of-sample predictive R^2 coefficient:

$$R^2_{OOS} = 1 - \frac{\sum_{(i,t)\in T_3}\left(r_{i,t+1} - \hat{r}_{i,t+1}\right)^2}{\sum_{(i,t)\in T_3} {r_{i,t+1}}^2}, \tag{4}$$

where $\hat{r}_{i,t+1}$ and $r_{i,t+1}$ indicated forecasted and realized weekly returns on the cryptocurrency i in week t+1, and T_3 denotes the testing sample. As in Gu et al. (2020), we estimated R^2_{OOS} using the entire sample of all weekly return observations pooled across cryptocurrencies and time.

R^2_{OOS} is the most common measure of prediction performance; however, it might be problematic to interpret or could even be irrelevant in some cases. In particular, it may fail to reflect the perspective of a quantitative portfolio manager who sorts securities to form a portfolio. What matters for them would be how effectively a model ranks assets in line with their actual ex-post returns. In other words, can it distinguish between losers and winners? Nevertheless, the link between predicted and realized returns might be blurred for R^2_{OOS}, as the cross-sectional correlation becomes drowned in the return variances (Coqueret, 2022). As a result, investors may still realize measurable economic gains—even if R^2_{OOS} is negative (Kelly et al., 2022).

We also calculate simple average cross-sectional correlation coefficients to deal with the issues. To be specific, each week, we compute the Spearman rank-based and Pearson product momentum correlation coefficient between the expected and realized payoffs. Subsequently, we calculate their averages through the time-series dimension ($\bar{\rho}_S$, $\bar{\rho}_P$). The obtained values provide an intuitive snapshot of the relationship between the model expectations and realized profits.

3. Baseline Empirical Findings

This section summarizes the prediction performance of machine learning methods. First, we summarize the overall measures of forecasting accuracy. Next, we explore the contribution of individual cryptocurrency characteristics.

3.1. Prediction Performance

Table 3 presents the prediction performance for different machine learning methods. Observe first the R^2_{OOS} measures. The overall predictability level qualitatively resembles those documented in earlier asset class studies. With the top models generating R^2_{OOS} of approximately 0.7-0.8%, the forecasting accuracy for cryptocurrencies is stronger than for the U.S. stocks; meanwhile, Gu et al. (2022) reported values not exceeding 0.4%—even for the best methods. On the other hand, the cryptocurrency return predictability lags behind corporate bonds or emerging market equities—where R^2_{OOS} may reach 2% to 5% (Bali et al., 2022; Leippold et al., 2022).

[Insert Table 3 about here]

The OLS model—by far the simplest in our array—has positive R^2_{OOS} of 0.243%. Hence, it does not suffer severely from overfitting and allows for generating potentially helpful forecasts. This observation stands in contrast with the seminal study of Gu et al. (2020)—where OLS produced a sizeable negative R^2_{OOS} value. Moreover, since overfitting does not seem to pose a substantial problem, the models aimed at reducing the overfitting do not bring measurable performance improvement. In consequence, the dimension reduction techniques (PLS) and penalized regressions (LASSO, ENET), generate low—and even negative— R^2_{OOS}. Furthermore, SVM fails to outperform OLS measurably.

Yet, what really makes a difference, is accounting for nonlinearities and interactions in the cross-section of returns. Hence, the prediction accuracy of tree methods and neural networks thrive in our ranking. This observation emphasizes the crucial benefits of machine learning models: they allow one to capitalize on unique return patterns in big datasets, which cannot be captured with simple linear models. In consequence, R^2_{OOS} for GBRT, RF, NN1, and NN2 visibly exceed 0.7% in all cases. The best performing model is NN1, yield R^2_{OOS} of 0.826%. Notably, the outperformance of the single-layer network of the double-layer matches the tendency of "shallow" models to beat the "deep" ones; this is recognized not only in finance, but also in other files—such as bioinformatics or computer vision. Gu et al. (2020) attribute this phenomenon to the relative dearth of data and high noise-to-signal ratio in asset pricing applications.

Finally, Table 3 also uncovers the strong performance of the COMB method; this corroborates earlier observations of Cakici et al. (2022) from international stock markets. Forecast combinations diminish the model variance, reducing the forecast error. In consequence, its R^2_{OOS} is at par with tree models and neural networks—equaling 0.742%.

The bottom rows of Table 3 display the average weekly cross-sectional correlation coefficients. At first glance, their values reveal notable similarities with the R^2_{OOS} measure. For instance, the tree models generate solid performance. Likewise, COMB fares very well; in fact, $\bar{\rho}_S$ classifies it as the best method among all examined. Nevertheless, one notable difference stands out: the correlation measures consistently exceed zero for all models considered. In other words, despite low or even negative R^2_{OOS} scores, all the machine learning models can be forged into successful portfolios based on sorting. Overall, $\bar{\rho}_P$ ranges from 0.046 to 0.107 and $\bar{\rho}_S$ from 0.052 to 0.092—laying the foundations for profitable investment strategies.

Notably, even though the correlation levels may seem relatively low, they could be sufficient for producing substantial economic gains. This occurs due to idiosyncrasies in a portfolio setting being canceled out, which boosts the resulting predictability. Consequently, even the allegedly worst-performing methods—such as OLS—can still effectively distinguish future winners from losers.

3.2. Which Cryptocurrency Characteristics Matter?

One of the unique traits of machine learning methods is that they help to pinpoint which of the variables matter for future returns. Therefore, they allow introducing some order into the "factor zoo." Though the cryptocurrency asset pricing literature is not as rich as for stocks, yet, the number of documented cross-sectional returns already reaches into the dozens. Hence, we are interested in the relative importance of these signals.

To measure the contribution of specific variables to overall return predictability, we compute variable importance (VI) by following the approach developed by Kelly et al. (2019). To be precise, for each predictor, we calculate in R^2_{OOS} resulting from setting all variables to zero while keeping all else fixed. This allows us to capture the critical determinants of the cross-sectional variation in cryptocurrency returns while simultaneously accounting for the full battery of 34 signals in the system.

3.2.1. Individual Importance

Figure 2 illustrates the VI of the 34 characteristics considered in this study. First, most models largely agree on the critical drivers of cryptocurrency returns. The top five variables include idiosyncratic risk (*ivol*), CAPM alpha (*alpha*), maximum monthly return (*max*), nominal price (*prc*), and value at risk (var). Moreover, the linear models—such as OLS, PLS, and penalized regressions—emphasize the role of the distance to the 90-day high (*90dh*). On the other hand, neural networks accentuate the importance of various liquidity indicators—including illiquidity ratio (*illiq*), trading volume (*volume*), or turnover (*turn*). The discrepancy in the importance of these features in linear and nonlinear models indicates that their importance emanates mainly from market interactions or nonlinearities, which simple OLS or LASSO cannot capture.

[Insert Figure 2 about here]

Interestingly, the top variables in Figure 2 do not include some well-known predictors of cryptocurrency returns (such as size) or the simple momentum indicators (such as *r31_2*). Possibly, their role may be subsumed by more sophisticated variables—capturing more effectively the same economic phenomenon. For example, the simple momentum may be subsumed CAPM alpha (*alpha*) or the distance to the 90-day high—as suggested by Jia et al. (2022).

Last, it is worth noting that the top predictors are relatively simple. They are derived from purely market data—such as prices, volumes, and returns. On the other hand, for example, the on-chain measures play a minor role. This conclusion is gracious for the practitioners: one can form successful predictions for the cryptocurrency markets without access to sophisticated data, such as the on-chain measures. What matters most can be found in simple price and return data, which is much more readily available.

Figure 3 zooms further on the 10 most important variables and illustrates their relative importance. Specifically, we compare their VI across different methods. Interestingly, for many methods, most predictability is derived from a handful of key variables. This is particularly striking for regularized regressions, where the top variable – *ivol* – matters more than several subsequent predictors taken together. A similar pattern is also noticeable for other linear models, as well as the tree methods, where a couple of key features (such as *90dh*, *prc*, and *ivol*) matter incomparably more than others. The neural networks, in turn, seem more democratic; the VI appears more equally distributed among the predictors.

[Insert Figure 3 about here]

Figure 4 sheds light on the model sparsity from a different angle. Specifically, we report the total VI for the top variables aggregated using the method of Bali et al. (2022). To this end, we first rescale all VI for each model so that their sum equals one. Then, we aggregate the scores for the top three, five, and 10 predictors.

[Insert Figure 4 about here]

The analyses seen in Figure 4 corroborate our earlier conclusions. The regularized regression and tree extract their information from a small number of top variables. For LASSO and ENET, the aggregate importance of the three top covariates equals about 40%, and the top 10 account for almost 70% of the predictability. In turn, while the aggregate VI of the top 3 predictors is smaller for the tree models, that of the top 10 exceeds even 75%.

Contrary to the trees and regularized regressions, the neural networks are more democratic. The aggregate VI of the top three variables does not pass even 15%, and the top 10 barely crosses 40%. This implies that neural networks extract information from a broader and more balanced panel of variables than other methods. Overall, the results in Figure 4 emphasize the remarkable dissimilarities behind the mechanics of machine learning models.

3.2.2. Category Importance

Last, we are also interested in the contribution of different types of characteristics to overall predictability. Hence, we closely follow the approach of Bali et al. (2022) to estimate aggregate VI (rescaled previously to sum to 1) within the categories representing similar economic intuitions—as specified in Table 1. We aim to determine which of these groups matters the most for the cross-section of cryptocurrency returns. Figure 5 displays the results of this exercise.

[Insert Figure 5 about here]

The three categories that drive the prices the most are associated with past returns, volatility, and liquidity. For example, in the COMB model (which averages the forecasts of other models), they account for 27%, 23%, and 20% of the aggregate VI—respectively. The other measures appear less critical, with a more modest role played by the on-chain metrics. However, the category importance is comparably uneven and may differ substantially across the models. For example, the liquidity characteristics, which account for only 17% of the aggregate VI for LASSO and ENET models, are simultaneously the most crucial set for NN1—representing 39% of the total VI. Likewise, the volatility category contribution ranges from 12% (SVM) to 35% (LASSO, ENET). This observation once again emphasizes the structural differences between various machine learning algorithms.

4. Machine Learning Portfolios

Having established the basic predictive properties of the machine learning signals, we now turn to portfolio analysis. We investigate whether machine learning forecasts can be transformed into effective investment strategies. We consider various study periods and investment horizons.

4.1. Portfolio Construction and Evaluation

Following the standard approach in asset pricing literature, we form one-way sorted portfolios. Specifically, each week, we rank all assets on their expected returns as predicted by different machine learning models. Next, we sort them into quintiles to form equal- and value-weighted portfolios.[4] In addition, we create long-short strategies that buy (sell) the quintile of cryptocurrencies with the highest (lowest) predicted return.[5] The performance of this portfolio provides an intuitive snapshot of the cross-sectional patterns of cryptocurrency returns.

We evaluate the portfolio performance with three distinct asset pricing models with factors derived from the cryptocurrency space. To begin with, we employ a simple one-factor market model:

$$r_t = \alpha_1 + \beta_{MKT} MKT_t + \varepsilon_t, \tag{5}$$

where r_t is the excess return on an examined cryptocurrency portfolio in week t; MKT_t is the market risk factor returns, ε_t is the residual term, β_{MKT} measures the market factor exposure, and α_1 is the weekly abnormal return (the so-called Jensen's alpha). The MKT return is calculated as a value-weighted excess return on all cryptocurrencies in the sample.

[4] The size distribution of the cryptocurrencies is remarkably uneven, with a handful of biggest assets accounting for most of the total market capitalization. Hence, to form more balanced yet tradeable portfolios, we cap the weight of the largest cryptocurrencies in the value-weighted portfolios. Specifically, we adopt the procedure from Jensen et al. (2022) and winsorize the market capitalization at the 90th percentile each week.

[5] As noted by Bianchi and Babiak (2022), though short selling cryptocurrencies is not always straightforward, it can be achieved on several major exchanges—such as Binance, Bitfinex, or Poloniex. The implementation requires borrowing assets at the current market price, selling them in the market, and repurchasing later to cover the position. Alternatively, the short positions in the hedge strategies could be interpreted as relative underweighting of a relevant market benchmark.

The second model, advocated by Liu et al. (2022), extends model (5) with two additional factors that represent size and momentum effects:

$$r_t = \alpha_3 + \beta_{MKT} MKT_t + \beta_{SIZE} SIZE_t^{LTW} + \beta_{MOM} MOM_t^{LTW} + \varepsilon_t. \quad (6)$$

The calculation of factor returns for model (6) reproduces the procedures in Liu et al. (2022, p. 1150). Specifically, the $SIZE_t^{LTW}$ factor is represented by a long-short portfolio that buys (sells) 30% of cryptocurrencies with the lowest (highest) market capitalization. On the other hand, the MOM_t^{LTW} factor comes from 2×3 sorts on market value (size) and momentum (r22_1). First, we split the market into big and small cryptos by their median. Second, we determine the 30th and 70th *r22_1* percentiles. Next, we intersect the two sorts to obtain six double-sorted portfolios. The momentum factor is then computed as the average return on two high-momentum portfolios minus the average return on the two low-momentum portfolios. The superscripts LTW highlight that the factors are calculated using the Liu et al. (2022) methodology.

Last, we extend the three-factor model (6) with three further factors considered in Babiak and Bianchi (2022):

$$r_t = \alpha_6 + \beta_{MKT} MKT_t + \beta_{SIZE} SIZE_t^{BB} + \beta_{MOM} MOM_t^{BB} + \beta_{LIQ} LIQ_t^{BB} + \beta_{VOL} VOL_t^{BB} + \beta_{REV} REV_t^{BB} + \varepsilon_t. \quad (7)$$

The model encompasses six observable risk factors included in the static F6 model from Babiak and Bianchi (2022); we then calculate them following the methods therein. Three new factors—LIQ_t^{BB}, VOL_t^{BB}, and REV_t^{BB}– represent the liquidity (*illiq*), volatility (*rvol*), and short-term reversal effects (*r2_1*). All the cross-sectional factors in the model are represented by long-short portfolios buying (selling) a quintile of assets with the highest (lowest) expected return.

Table A1 in the Online Appendix summarizes the statistical properties of returns on factor portfolios from different asset pricing models. In line with earlier studies (Liu et al., 2022; Babiak & Bianchi, 2022), all factors generate positive average returns in our sample in the range of 0.62% (MKT) to 6.49% (LIQ). Their level qualitatively resembles evidence from the referenced studies. Notably, all factors exhibit substantial volatility—with the return standard deviation from 6.24% to 15.48% per week.

4.2. Baseline Performance

Table 4 reports the average weekly returns on different machine learning strategies, with Panels A and B focusing on equal- and value-weighted portfolios—respectively. A quick look at the results uncovers a crucial pattern: all machine learning forecasts can be transformed into measurable economic gains for investors. All prediction models can effectively sort cryptocurrencies in line with their realized returns, so quintile portfolios systematically exhibit a monotonic or near-monotonic pattern. In consequence, all long-short strategies produce robust profits—even those that display seemingly low R^2_{OOS} values (such as LASSO and ENET).

[Insert Table 4 about here]

Consider first the equal-weighted portfolios (Panel A). The mean weekly returns span from 2.10% (PLS) to 4.32% (COMB), all statistically significant. The best-performing strategy is the forecast combination, highlighting the benefits of averaging predictions of several models. The superior profits on the COMB portfolios corroborate similar findings of Cakici et al. (2022) for international stock markets, all strategies have their benefits; however, combining them proves particularly efficient.

Besides COMB, the highest returns are earned by the GBRT and RF portfolios—closely followed by NN1 and NN2. The outperformance of the tree models and neural networks emphasizes the essential role of nonlinearities and interactions in cryptocurrency markets. The linear models, such as raw and regularized regressions or dimension reduction techniques, cannot capture these inherent features of the return cross-section. Consequently, their mean returns—though still highly significant—visibly lag.

The machine learning profits strategies cannot be subsumed by popular risk factors. The long-short portfolios continue to generate significant alphas after accounting for the factor models (5) to (7). Even the most comprehensive six-factor model fails to explain the machine learning returns. In consequence, for example, α_6 on the long-short COMB portfolio equals 3.96% (t-stat = 11.88). To summarize, the machine learning models extract new information from cryptocurrency markets that cannot be easily captured with standard linear models.

Panel B of Table 4 reports the returns on value-weighted portfolios. The results uncover essentially the same patterns. All the long-short portfolios generate positive and highly significant profits, which any of our factor models cannot capture. The raw and abnormal returns magnitude is qualitatively similar to the equal-weighted strategies. For example, the three-factor alphas of Liu et al. (2022) on the value-weighted portfolios range from 2.11% to 4.73%; meanwhile, for the equal-weighted strategies, they fare between 1.86% and 4.28%. In a nutshell, the machine learning effectiveness works independently of the portfolio weighting scheme. This observation may matter particularly from a practical angle, as it limits the reliance on overweighting small and illiquid coins.

While our discussion has so far revolved around the hedge portfolios, a closer look at the individual quintiles uncovers further insights. The cross-sectional return pattern suggests that the alphas on long-short strategies mainly originate from their short legs rather than long ones. The mean returns on the cryptocurrency quintile with the highest expected returns are relatively moderate. For instance, for the value-weighted portfolios (Table 4, Panel B), they range from 0.14% (PLS) to 1.00% (NN1). On the other hand, the negative returns on the bottom portfolios are markedly higher in absolute terms: from -3.84% (COMB) to -2.05% (PLS).

To illustrate this further, Figure 6 displays alphas from the model of Liu et al. (2022). The cross-sectional pattern in cryptocurrency returns is somewhat L-shaped rather than linear—especially for the linear models. The low abnormal returns on the bottom quintile stand out, surpassing all other portfolios in absolute terms.

[Insert Figure 6 about here]

Substantially higher alphas on the bottom quintiles comply with the mispricing view on market anomalies (Stambaugh et al., 2012, 2015; Avramov et al., 2013, 2019). Buying—stocks and cryptocurrencies alike—is easier than shorting for most investors. In consequence, arbitrageurs may be quicker to eliminate underpricing than overpricing. In the extreme, when the short selling is severely impeded, the overpricing may last over a prolonged period. This arbitrage asymmetry may result in stronger abnormal returns on the short legs of anomaly portfolios than on long ones.

Notably, the reasoning above inherently links the cross-sectional return patterns with mispricing. Hence, the strong alphas on the short legs of our portfolios may suggest a similar economic mechanism: the return predictability in cryptocurrency markets is due to mispricing rather than risk.

4.3. Subperiod Analysis

The return predictability in financial markets declines over time. A voluminous literature has documented that the profitability of stock characteristics has been declining due to investor learning, improvements in liquidity, or an increase in arbitrage activity (Schwert, 2003; Chordia et al., 2014; McLean & Pontiff, 2016; Caluzzo et al., 2019). Furthermore, this phenomenon extends beyond mere stocks: to currencies, industry and country portfolios, or even mutual funds (Bartram et al., 2018; Zaremba et al., 2020; Jones & Mo, 2021).[6] This decrease in predictability typically connects with the mispricing story of asset pricing anomalies: once the market becomes more efficient, mispricing becomes less common.

Our findings from Section 4.2 concerning the alphas in the long and short legs align with the mispricing story. Hence, do the abnormal returns share the same time-series patterns

[6] It is important to note that Jacobs and Müller (2020) do not confirm a robust predictability decline in international stock markets.

as in the stock market landscape? Do they gradually decline over time as cryptocurrency markets become more liquid and efficient?

Figure 7 illustrates the cumulative returns on the long-short machine learning strategies that are considered in this study. To begin with, their performance is remarkably stable over time. No major downturns occurred within our study period, regardless of the situation in the cryptocurrency markets. The long-run fluctuations in returns are limited.

[Insert Figure 7 about here]

Despite this apparent stability, a particular pattern seems noticeable: the long-short portfolio returns have declined over time. In particular, the years 2021 and 2022 witnessed a flattening of the profit curve. While all strategies seem to continue generating abnormal returns, their size appears lower.

To scrutinize this effect further, we reproduce our portfolio sorts from Table 4 in subperiods. Specifically, we split the whole research period into rough halve (January 17, 2018 to April 7, 2020 and April 8, 2020 to July 6, 2022) and check the performance therein. Table 5 reports the results of this exercise. For brevity, we limit the presentation to value-weighted strategies.[7]

[Insert Table 5 about here]

Both subperiods exhibit substantial return predictability. The mean returns and alphas are positive and highly significant for all models in both halves of the research sample. Nevertheless, the disparity in performance in the earlier and later years is evident. The mean long-short portfolio returns from January 2018 to April 2020 varied from 2.85% (PLS) to 5.49% (COMB). The latter years, however, brought a measurable drop in profitability. The mean returns declined to a range between 1.57% (PLS) to 3.96% (COMB). On average, the mean long-short strategy payoffs decreased by 30% and were accompanied by a similar reduction in alphas. Notably, this decrease occurred despite the growing number of listed assets—which generally tend to boost profits on hedge portfolios (Bessembinder et al., 2021).

To conclude, our findings from the cryptocurrency markets match earlier evidence from stocks and other asset classes. Whether it is due to improved liquidity, market efficiency, or investor learning, abnormal returns' magnitude gradually declines over time. This finding complies with the mispricing narrative of return predictability.

[7] The results for equal-weighted portfolios are available upon request.

4.4. Alternative Forecast Horizons

Our tests so far have concentrated on a one-week prediction horizon. However, return predictability with machine learning may vary substantially across different horizons. For example, Coqueret (2022) documents that predictability is typically stronger for longer forecast horizons. Consistently with this, Gu et al. (2022) and Leippold et al. (2022) observe higher R^2_{OOS} for longer horizons. The effect may be associated with the higher signal-to-noise ratio in long-term returns and more pronounced signal persistence. Does the prediction horizon also matter for machine learning forecasts in cryptocurrencies?

To answer this question, we reproduce our baseline analysis from Section 4.2 using two alternative forecast horizons: a shorter one (one day) and a longer one (two weeks). We begin by assessing the prediction accuracy. Holding all else equal, we recalculate the performance measures from Table 3 for these two alternative settings. Table 6, Panel A uncovers the results.

[Insert Table 6 about here]

The R^2_{OOS} values and average correlation coefficients for daily forecasts (Panel A.1) are considerably lower than for the weekly intervals. For example, the average Pearson measure for the COMB model equals 0.022 versus 0.096 in our baseline approach. All metrics and modes exhibit a similar drop in the obtained scores. On the other hand, the predictions of the biweekly returns (Panel A.2) are slightly higher than in our baseline case. For instance, the same Pearson measure for COMB—evoked in the example above—amounts to 0.139. Furthermore, the R^2_{OOS} measures exhibit a noticeable increase and the effect holds across all the prediction models. In other words, as in the stock markets, predictability tends to be weaker for shorter horizons and vice versa.

Can this boost in predictability be forged into higher profits on the machine learning strategies? To verify this, we revisit our univariate portfolio sorts from Section 4.2 with a modified prediction horizon and holding period. Specifically, we use signals from the daily prediction models to form long-short strategies that are reformed daily. Analogously, we also employ the 14-day predictions to form bi-weekly reconstructed portfolios. Table 6, Panel B reports the returns on long-short value-weighted machine learning portfolios implemented using the aforementioned procedures.

In line with the observations in Panel A, the strategies based on a daily horizon display noticeably lower profitability. The mean returns and alphas decline substantially compared to our baseline-weekly horizon, though all remain highly significant. The mean weekly returns span between 0.34% (SVM) and 2.95% (COMB). Conversely, the bi-weekly strategies show marginal improvement relative to our baseline case; however, the change is minuscule. For example, the COMB alpha equals 4.89% versus 4.73% in the baseline scenario. Moreover, the increase in profitability is uneven, and some models—such as GBRT and RF—witness even a slight decrease in abnormal returns.

To sum up, the forecast horizon indeed affects machine learning performance. In particular, reducing the return interval to one day shrinks the prediction accuracy and portfolio performance measurably. In turn, an increase in the return horizon plays a smaller role; this results in only in a minor improvement in prediction accuracy and a negligible boost in portfolio profits.

5. Return Predictability and Limits to Arbitrage

Machine learning methods aggregate multiple signals from individual anomalies. Accordingly, any weaknesses of these signals may transmit into the models' forecasts. This inference may have substantial consequences for machine learning effectiveness. For example, in the equity universe, most well-known anomalies—such as momentum, value, or profitability—extract their alphas mainly from small and illiquid firms (Hong et al., 2000; Israel & Moskowitz, 2013; Hou et al., 2020; Cakici & Zaremba, 2022). This evidence aligns with the view that mispricing typically arises where arbitrage activities are constrained. In line with this, as shown by Avramov et al. (2022), machine learning profits also tend to concentrate on difficult-to-arbitrage stocks.

Our evidence so far complies with the mispricing view on return predictability. Hence, if the cryptocurrency return predictability is driven by mispricing, it should prevail in difficult-to-arbitrage assets. So, does it? We run two experiments to shed light on these issues: cross-sectional regressions with interactions and bivariate portfolio sorts.

5.1. Cross-Sectional Regressions With Interaction Terms

To identify the difficult-to-trade securities, we employ three popular proxies of limits of arbitrage: idiosyncratic volatility (*ivol*), illiquidity (*illiq*), and bid-ask spread (*bidask*). Numerous asset pricing studies have argued that assets with values of these variables are costly to trade and difficult to hedge; consequently, any mispricing is likely to persist (e.g., Pontiff, 2006, 2016; Sadka & Scherbina, 2007; Chordia et al., 2008; McLean, 2010; Lam & Wei, 2011). In addition to these three variables, we also calculate an aggregate measure of limits of arbitrage (*lim*). For each cryptocurrency, it is represented by its average cross-sectional rank on *ivol*, *illiq*, and *bidask* in a given week.

To explore the impact of the limits to arbitrage on return predictability, we estimate weekly cross-sectional regressions with interaction terms accounting for the limits to arbitrage:

$$r_{i,t+1} = \gamma_{0,t} + \gamma_{1,t}E_t(r_{i,t+1}) + \gamma_{2,t}Arb_{i,t} + \gamma_{3,t}E_t(r_{i,t+1})Arb_{i,t} + \varepsilon_{i,t+1}, \quad (8)$$

where $E_t(r_{i,t+1})$ is the cryptocurrency i expected return, as predicted by the COMB model, and $r_{i,t+1}$ is its realized value. For brevity, we limit our presentation to the COMB predictions, which aggregates all individual models; using other models yields qualitative results and is available upon request. $Arb_{i,t}$ is a dummy variable taking a value of one if a given measure of limits to arbitrage (*ivol*, *illiq*, *bidask*, *lim*) is higher than the weekly cross-sectional median or zero otherwise. $\varepsilon_{i,t+1}$ is the error term and $\gamma_{0,t}$, $\gamma_{1,t}$, $\gamma_{2,t}$, and $\gamma_{3,t}$ are the estimated weekly regression coefficients. Our focus is on the coefficient $\gamma_{3,t}$; its positive value would imply that high limits to arbitrage augment the return predictability by machine learning models.

Table 7 reports the average slope coefficients from regression (8). The limits to arbitrage clearly matter. The $\gamma_{3,t}$ coefficients are positive and highly significant in all specifications, regardless of the particular proxy used. Admittedly, the coefficients on $E_t(r_{i,t+1})$ also remain significant in all cases; this signifies that machine learning models continue to forecast returns successfully even if arbitrage constraints are controlled for. In other words, they work even in market segments where trading is relatively more straightforward. Nonetheless, the return predictability in difficult-to-trade assets is measurably stronger. This observation lends further support to the hypothesis that return predictability in the cryptocurrency markets originates mainly from mispricing.

[Insert Table 7 about here]

5.2. Bivariate Portfolio Sorts

We carry on with bivariate portfolio sorts to corroborate the conclusions from the cross-sectional regressions seen in Section 5.1. In this analysis—each week—we rank cryptocurrencies on one of the four proxies for arbitrage constraints, *ivol*, *illiq*, *bidask*, and *lim*, and sort them into terciles. Thus, we generate three market segments with high, medium, and low arbitrage difficulties. Next, within each of these segments, we sort the assets again on their predicted returns from a machine learning model. Again, for conciseness, we limit the presentation to the COMB model.[8] Next, we weight the cryptocurrencies—equally or on their capitalization—in order to obtain nine portfolios from double-dependent sorts. Finally, we calculate long-short machine learning strategies within each of the terciles representing different levels of limits to arbitrage. Table 8 presents the outcomes of this analysis.

[Insert Table 8 here]

[8] An application of other models leads to consistent conclusions. The detailed results are available upon request.

Table 8, Panel A reports the average portfolio returns. Overall, the machine learning strategies work effectively across all market segments. The long-short portfolios generate positive and significant mean raw and abnormal returns in each of the considered specifications. Nonetheless, their magnitude evidently differs. The payoffs are considerably more solid in the market segments with high barriers to arbitrage than in those when trading is more straightforward.

Take the aggregate *lim* proxy for arbitrage constraints as an example (Table 8, Panel A.4). The value-weighted long-short strategy (right section of Panel A) earns an average return of 5.53% (t-stat = 6.06) per week in the high *lim* tercile. On the other hand, in the low arbitrage segment, the equivalent mean return amounts to 1.96% (t-stat = 5.15). The difference between the high and low *lim* segments equals 3.57% and is highly significant (t-stat = 3.76). Furthermore, it cannot be explained by common risk factors from the model of Liu et al. (2022). The equal-weighted portfolios yield similar results, as seen in the left section of Panel A.

Panels A.1 to A.3 of Table 8 shed further light on the role of different components of the *lim* measure. Which proxies for the limits to arbitrage matter the most? The strongest performance dispersion is observed for *ivol*. Consider the value-weighted portfolios. In this case, the average difference between long-short portfolio alpha in the high and low *ivol* terciles equals 5.38% (t-stat 6.69). On the other hand, for *bidask*, the effect is relatively weaker—though still visible. The alpha difference between high and low *bidask* tercile amounts to 1.88% (t-stat = 2.08).

Our earlier evidence in Section 4.2 indicates that most of the abnormal returns originate from the short legs of investment strategies. A closer look at individual bivariate portfolios in Table 8, Panel A leads to similar conclusions. However, the absolute value of returns is not equal across different market segments. Concretely, it is powerful in terciles with high limits to arbitrate and relatively weaker in the terciles with low constraints. Consider the two-way sorts on *lim.* The average returns on the cryptocurrencies with the highest expected returns span between 0.24% and 0.85%; on the other hand, the ones with the lowest predictions are between -5.29% and -1.30%. Notably, the meager mean return of -5.29% is recorded precisely on the portfolio with the highest *lim* values. In other words, the machine learning strategies derive their predictability—particularly from short positions in difficult-to-arbitrage securities.

The conclusions above are especially interesting when juxtaposed with the average cryptocurrency capitalization in different market segments and their respective relative size. Table 8, Panel B uncovers the details of such distribution. Consider the bivariate sorts on expected returns in *lim* in Panel B.4. The tercile of cryptocurrencies with the lowest limits to arbitrage (the one when machine learning profits are the weakest) accounts for 96.45% of the cryptocurrency market capitalization. Meanwhile, the tercile of the most difficult-to-trade coins (the ones which generate the highest profits) represent only 0.80% of the market value. In other words, the market segment where the return predictability is the most robust accounts for a minuscule fraction of the cryptocurrency universe.

Furthermore, the two-way sorted portfolio with the lowest expected return in the high *lim* tercile—which produced a remarkable mean value-weighted return in Panel B of –5.29%—represents only 0.08% of the market value, and the average cryptocurrency capitalization therein equals $ 20 million. In consequence, shorting and effective trading of such small cryptocurrencies may be challenging. Moreover, their economic significance is rather modest.

To conclude, the machine learning models work across diverse segments of cryptocurrency markets. Nevertheless, the return predictability is the strongest where trading is the most difficult. Particularly, it originates from short positions in difficult-to-trade assets—which aggregate market value; because of this, economic importance may be limited.

6. Practical Investor Perspective

Finally, we are also interested in the practical perspective on return predictability with machine learning. Do these forecasts translate into feasible strategies? Can they survive real-world conditions? To throw light on these issues, we look closer at three aspects: (i) risk statistics; (ii) portfolio investability; and (iii) trading costs.

6.1. Risk Statistics

Table 9, Panel A reports the major risk measures for the long-short value-weighted machine learning strategies. Their standard deviations of weekly returns fit into a moderate range of between 5.58% (OLS) and 6.65% (SVM). Consequently, the strategies display impressive Sharpe ratios—their annualized values fall between 2.73 (PLS) and 5.37 (COMB). Notably, this level of risk-adjusted returns noticeably beats their counterparts in traditional asset classes—such as equities. For example, Gu et al. (2020) report annualized Sharpe ratios on long-short machine learning strategies for the U.S. stocks of up to 2.45; furthermore, Leippoled et al. (2022) record values of up to 3.45 for the Chinese market.

[Insert Table 9 about here]

Notably, the long-short cryptocurrency strategies also exhibit asymmetrical risks—which are not fully captured by the standard deviation measure. For example, maximum weekly losses on all strategies are substantial; in many cases, an investor must be prepared to lose more than 20% in one week. The worst performer in this regard is SVM, with a maximum seven-day loss of 33.43%. This tail risk is also manifested in prolonged and deep maximum drawdowns. In most cases, the investors would have to survive a loss of more than 30-40% of the invested capital. For SVM, this statistic amounts to as much as 58.83%. Hence, the potential tail risk in cryptocurrency machine-learning strategies may be substantial.

6.2. Long-Short Portfolio Feasibility

Table 9, Panel B focuses on the average cryptocurrency capitalization of different quintiles. The top quintile typically contains bigger coins, with an average market value between three and nine billion USD. Nevertheless, as seen in both Table 4 and Figure 6, the abnormal returns concentrate mainly in the short leg; i.e., the one with the lowest expected return. In this case, however, the capitalizations are noticeably smaller. The average cryptocurrency market value in the bottom quintile is roughly 100 to 600 million USD.

Table 9, Panel C offers an additional perspective on this issue by reporting the average proportion of the total market capitalization captured by different quintiles. The long legs typically account for the biggest chunk of the cryptocurrency universe, representing between 26.3% (NN1) and 77.4% (LASSO) of its aggregate market value. On the other hand, the bottom quintile—the primary source of alphas in the long-short strategies—accounts for between 0.8% (COMB) and 9.1% (NN1) of the market. Notably, this subset does not capture more than 1% of the aggregate cryptocurrency capitalization for almost half the strategies.

To sum up, a practical implementation of long-short cryptocurrency strategies may pose substantial challenges. It hangs on taking massive short positions in very small coins that represent a minuscule fraction of the entire market. The feasibility of such operations may be questionable. The short-selling opportunities may be limited due to lending fees, deposit requirements, or unavailability of assets for lending.

6.3. Trading Costs

Machine learning strategies typically require active trading and frequent portfolio reconstructions (Gu et al., 2020; Leippold et al., 2022). This, in turn, may incur substantial trading costs—especially since the strategies derived their profitability primarily from small coins. High portfolio turnover and trading costs are the enemies of many quantitative strategies, impeding their performance and even calling into question their validity (Novy-Marx & Velikov, 2016). Does this problem haunt cryptocurrency strategies as well?

6.3.1. Portfolio Turnover

Table 9, Panel D demonstrates the portfolio turnover of quintile portfolios formed on machine learning forecasts. Analogously, as in Bollerslev et al. (2018) and Koijen (2018), we calculate the turnover statistic as the average of portfolio share that needs to be replaced each week:

$$PT_t = \frac{1}{2}\sum_{i=1}^{n}\left|w_{i,t-1} \times \left(1 + r_{i,t}\right) - w_{i,t}\right|, \tag{9}$$

where $w_{i,t-1}$ and $w_{i,t}$ are the weights of cryptocurrency i in two consecutive weeks, and $r_{i,t}$ is the crypto return. We compute a one-sided (rather than two-sided) turnover measure to avoid double-counting of buys and sells.[9]

Consistent with earlier machine learning studies, the portfolio turnover is substantial. The individual quintile portfolios typically require replacing about 40 to 70% of their value each week. As a result, the long-short strategies exhibit a weekly turnover that falls between 78.6% (LASSO) and 110.8% (SVM). The approach based on the COMB model, previously identified as the top performer, has a turnover ratio of 83.9%. Such a substantial portfolio rotation may negatively impact real-world portfolio performance.

6.3.2. Trading Cost Estimation

To quantify the transaction cost drag on the profitability of machine learning strategies, we continue with a more formal analysis. To ensure a comprehensive perspective, we calculate three different sensitivity measures to trading costs. First, we compute the breakeven point for the trading costs. We estimate it as the average portfolio return divided by the average turnover.

Second, we turn to an estimation of net returns. In practice, obtaining detailed data on a large sample of cryptocurrency daily spreads and slippage rates may be challenging. Furthermore, the fixed transaction fees may vary across exchanges and investor types; they are also subject to various special offers. Hence, in the most simplistic approach, we follow Bianchi et al. (2022) and assume a fixed one-way fee of 30 (40) basis points for the long (short) leg as an all-inclusive approximation of transaction costs. Bianchi et al. (2022, p. 6) deem this measure as “a set of fairly conservative transaction costs for a market maker.”

A vital shortcoming of the fixed trading costs is that they miss the inherent features of the machine learning strategies, which derive profits mainly from small cryptocurrencies. These coins are more likely to be illiquid and difficult to trade; and so the overall

---

[9] Notably, the average weekly turnover computed this way for the case of long-short strategies may exceed 100% as the long and short legs’ turnover are added up.

transaction costs may be underestimated. Hence, we also employ a third, more nuanced measure: the return net of variable transaction costs. In this framework, we dynamically estimate each cryptocurrency's effective spread using the methods of Corwin and Schultz (2012) and Abdi and Ranaldo (2017). This approach's total one-way transaction cost is calculated as a sum of the bid-ask spread estimate (*bidask* defined as in Table 1), with an additional flat fee equal to 10 basis points. As seen from the properties of *bidask* in Table 2, this measure is much more conservative and assumes higher trading costs—especially for the tail of less liquid coins.

6.3.3. Cost-Adjusted Returns

Table 10 summarizes the role of trading costs in the long-short machine learning strategies illustrated using different approaches. Since the results for equal-weighted (Panel A) and value-weighted portfolios (Panel B) are relatively similar, let us focus on the value-weighted strategies. The breakeven trading costs range between 133 (SVM) and 281 (COMB) basis points. These numbers are far beyond the conservative trading costs assumptions in Babiak et al. (2022); thus, they promise a potentially successful portfolio implementation. The combination strategy, which fared the best in Table 4, also turns out most resilient to the impact of trading costs.

[Insert Table 10 about here]

Not surprisingly, all strategies produce positive and significant mean returns when the fixed trading costs are considered. The average profits range from 1.66% (PLS) to 4.13% (COMB). More importantly, most strategies also survive the variable transaction cost approach. Admittedly, the drag on portfolio performance is substantial; moreover, the strategies lose about 60% of their gross profits on average. Furthermore, the mean returns on two long-short portfolios—PLS and SVM—are rendered insignificant. Nevertheless, most algorithms survive—continuing to produce sizeable and significant profits. Our top performer, the COMB strategy, generates an average weekly net return of 2.89% (*t*-stat = 5.54).

To sum up, the cryptocurrency machine learning strategies seem robust to real-world impediments. Despite a high portfolio turnover, they endure even conservative trading costs. Nevertheless, this performance comes with two caveats. First, the strategies exhibit considerable tail risk and large drawdowns. Second, it requires taking substantial short positions in the smallest cryptocurrencies—which pose practical difficulties. A successful practical implementation hinges on resolving these two challenges.

7. Concluding Remarks

Both machine learning tools and cryptocurrency investments have attracted growing investor attention in recent years. Our study intersects these two fields. We combine asset pricing research with machine learning to gain new insights into return predictability in cryptocurrency markets. Using data ranging from 2017 to 2022, we calculate 34 cryptocurrency features. We use 10 popular machine learning algorithms—including regularized regression, dimension reduction techniques, tree models, and neural networks. Using the cryptocurrency characteristics as inputs, we train the models to predict the cross-section of returns.

Our empirical analysis brings four main contributions. First, we demonstrate that machine learning models can be effectively applied to predict the cross-section of cryptocurrency returns. All our models' forecasts generate measurable economic gains. In consequence, they translate into profitable portfolios from univariate sorts. A value-weighted long-short quintile portfolio formed on the most successful model—the forecast combination—generates a mean monthly return of 4.72%. These abnormal returns cannot be subsumed by popular factor models and hold in various robustness checks, different weighting schemes, forecast horizons, and subperiod analyses.

Second, we cast light on the crucial determinants of the cross-sectional variation in cryptocurrency returns. A variable importance analysis allows us to pinpoint the crucial variables. Most models derive their predictability from a handful of relatively simple signals—such as idiosyncratic volatility, CAPM alpha, maximum daily return, nominal price, value at risk, and distance to a 90-day high. With respect to general categories of cryptocurrency features, what matters the most are past returns, volatility, and liquidity.

Third, we provide insights into the sources of the cross-sectional return predictability in cryptocurrency markets. Our results comply with the mispricing view on return predictability. The machine learning profits derive mainly from short legs, where eliminating mispricing might be more challenging due to short-selling constraints. Furthermore, the alphas concentrate in difficult-to-trade assets—where limits to arbitrage are more pronounced. Finally, abnormal returns decline over time as markets become more efficient.

Fourth, we assess the cryptocurrency machine-learning strategies from a practical angle. Despite substantial portfolio turnover, we find them resilient to transaction costs—even when we assume relatively high trading costs. Nonetheless, the allegedly impressive performance has two hooks. First, it exhibits substantial tail risk that manifests in sizeable drawdowns. Second, it requires taking substantial short positions in the small coins. Resolving both these issues may be critical for a successful implementation of cryptocurrency machine-learning strategies.

References

Abdi, F., & Ranaldo, A. (2017). A simple estimation of bid-ask spreads from daily close, high, and low prices. *Review of Financial Studies*, 30(12), 4437-4480.

Alexander, C., & Dakos, M. (2020). A critical investigation of cryptocurrency data and analysis. *Quantitative Finance*, 20(2), 173-188.

Amihud, Y. (2002). Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets*, 5(1), 31-56.

Avramov, D., Cheng, S., & Metzker, L. (2021). Machine learning versus economic restrictions: Evidence from stock return predictability. *Management Science*, forthcoming.

Avramov, D., Chordia, T., Jostova, G., & Philipov, A. (2013). Anomalies and financial distress. *Journal of Financial Economics*, 108(1), 139-159.

Avramov, D., Chordia, T., Jostova, G., & Philipov, A. (2019). Bonds, stocks, and sources of mispricing. George Mason University School of Business Research paper, (18-5).

Azevedo, V., Kaiser, S., & Müller, S. (2022). Stock market anomalies and machine learning across the globe. Available at SSRN 4071852.

Bali, T. G., Cakici, N., & Whitelaw, R. F. (2011). Maxing out: Stocks as lotteries and the cross-section of expected returns. *Journal of Financial Economics*, *99*(2), 427-446.

Bali, T. G., Goyal, A., Huang, D., Jiang, F., & Wen, Q. (2022). Predicting corporate bond returns: Merton meets machine learning. Georgetown McDonough School of Business Research Paper (3686164), 20-110.

Bartram, S.M., Djuranovik, L., & Garratt, A. 2018. Currency Anomalies. 31st Australasian Finance and Banking Conference 2018, Available at SSRN: https://ssrn.com/abstract=3222252 or http://dx.doi.org/10.2139/ssrn.3222252.

Bates, J. M., & Granger, C. W. (1969). The combination of forecasts. *Journal of the Operational Research Society*, *20*(4), 451-468.

Bessembinder, H., Burt, A.P., & Hrdlicka, C.M. (2021). Time series variation in the factor zoo. Available at SSRN: https://ssrn.com/abstract=3992041 or http://dx.doi.org/10.2139/ssrn.3992041.

Bhambhwani, S., Delikouras, S., & Korniotis, G. M. (2019). Do fundamentals drive cryptocurrency prices? CEPR Discussion Papers 13724, Centre for Economic Policy Research.

Bianchi, D., & Babiak, M. (2022). A factor model for cryptocurrency returns. Available at SSRN 3935934.

Bianchi, D., Babiak, M., & Dickerson, A. (2022). Trading volume and liquidity provision in cryptocurrency markets. *Journal of Banking & Finance*, 106547.

Bianchi, D., Büchner, M., & Tamoni, A. (2021). Bond risk premiums with machine learning. *Review of Financial Studies*, 34(2), 1046-1089.

Borri, N. (2019). Conditional tail-risk in cryptocurrency markets. *Journal of Empirical Finance*, *50*, 1-19.

Borri, N., & Shakhnov, K. (2022). The cross-section of cryptocurrency returns. *Review of Asset Pricing Studies*, 12(3), 667-705.

Borri, N., Massacci, D., Rubin, M., & Ruzzi, D. (2022). Crypto risk premia. Available at SSRN 4154627.

Bouri, E., Kristoufek, L., Ahmad, T., & Shahzad, S. J. H. (2022). Microstructure noise and idiosyncratic volatility anomalies in cryptocurrencies. *Annals of Operations Research*, 1-27.

Brauneis, A., R. Mestel, Riordan, R., & Theissen, E. (2021). How to measure the liquidity of cryptocurrency markets? *Journal of Banking & Finance*, 124, 106041.

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

Burggraf, T., & Rudolf, M. (2021). Cryptocurrencies and the low volatility anomaly. *Finance Research Letters*, 40, 101683.

Cai, C. X., & Zhao, R. (2021). Salience theory and cryptocurrency returns. *Available at SSRN 3983602*.

Cakici, N., & Zaremba, A. (2022). Salience theory and the cross-section of stock returns: International and further evidence. *Journal of Financial Economics*, 146(2), 689-725.

Cakici, N., Fieberg, C., Metko, D., & Zaremba, A. (2022). Machine learning goes global: Cross-sectional return predictability in international stock markets. Available at SSRN 4141663.

Calluzzo, P., Moneta, F., & Topaloglu, S. (2019). When anomalies are publicized broadly, do institutions trade accordingly? *Management Science*, 65(10), 4555-4574.

Calluzzo, P., Moneta, F., & Topaloglu, S. (2019). When anomalies are publicized broadly, do institutions trade accordingly? *Management Science*, 65(10), 4555-4574.

Chen, L., Pelger, M., & Zhu, J. (2019). Deep learning in asset pricing. arXiv preprint arXiv:1904.00745.

Chen, R., Lepori, G. M., Tai, C. C., & Sung, M. C. (2022). Can salience theory explain investor behaviour? Real-world evidence from the cryptocurrency market. *International Review of Financial Analysis*, 84, 102419.

Chordia, T., Roll, R., & Subrahmanyam, A. (2008). Liquidity and market efficiency. *Journal of Financial Economics*, 87(2), 249-268.

Chordia, T., Subrahmanyam, A., & Tong, Q. (2014). Have capital market anomalies attenuated in the recent era of high liquidity and trading activity? *Journal of Accounting and Economics*, 58(1), 41-58.

Clemen, R.T. (1989). Combining forecasts: A review and annotated bibliography. *International Journal of Forecasting*, *5*(4), 559–583.

Cong, L. W. & He, Z. (2019). Blockchain disruption and smart contracts. Re*view of Financial Studies*, 32(5), 1754-1797.

Cong, L. W., Karolyi, G. A., Tang, K., & Zhao, W. (2022). Value premium, network adoption, and factor pricing of crypto assets. Available SSRN 3985631.

Coqueret, G. (2022). Persistence in factor-based supervised learning models. *Journal of Finance and Data Science*, 8, 12-34.

Corwin, S. A., & Schultz, P. (2012). A simple way to estimate bid-ask spreads from daily high and low prices. *Journal of Finance*, 67(2), 719-760.

Cosemans, M., & Frehen, R. (2021). Salience theory and stock prices: Empirical evidence. *Journal of Financial Economics*, 140(2), 460-483.

Datar, V. T., Naik, N. Y., & Radcliffe, R. (1998). Liquidity and stock returns: An alternative test. *Journal of Financial Markets*, 1(2), 203-219.

Diebold, F. X., & Shin, M. (2019). Machine learning for regularized survey forecast combination: Partially-egalitarian LASSO and its derivatives. *International Journal of Forecasting*, *35*(4), 1679-1691.

Dobrynskaya, V. (2021). Cryptocurrency momentum and reversal. Available at SSRN 3913263.

Dobrynskaya, V., & Dubrovskiy, M. (2022). Cryptocurrencies meet equities: risk factors and asset pricing relationships. Higher School of Economics Research Paper No. WP BRP, 86.

Dong, B., Jiang, L., Liu, J., & Zhu, Y. (2022). Liquidity in the cryptocurrency market and commonalities across anomalies. *International Review of Financial Analysis*, *81*, 102097.

Drobetz, W., & Otto, T. (2021). Empirical asset pricing via machine learning: evidence from the European stock market. *Journal of Asset Management*, *22*(7), 507-538.

Fama, E. F., & French, K. R. (2008). Dissecting anomalies. *Journal of Finance*, *63*(4), 1653-1678.

Fama, E. F., & French, K. R. (2012). Size, value, and momentum in international stock returns. *Journal of Financial Economics*, *105*(3), 457-472.

Fama, E. F., & MacBeth, J. D. (1973). Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, 81(3), 607-636.

Fieberg, C., Metko, D., Poddig, T., & Loy, T. (2022). Machine learning techniques for cross-sectional equity returns' prediction. *OR Spectrum*, in press.

George, T. J., & Hwang, C. Y. (2004). The 52-week high and momentum investing. The *Journal of Finance*, 59(5), 2145-2176.

Giglio, S., Kelly, B., & Xiu, D. (2022). Factor models, machine learning, and asset pricing. *Annual Review of Financial Economics*, 14, 337-368.

Grobys, K., & Junttila, J. (2021). Speculation and lottery-like demand in cryptocurrency markets. *Journal of International Financial Markets, Institutions and Money*, *71*, 101289.

Grobys, K., & Sapkota, N. (2019). Cryptocurrencies and momentum. *Economics Letters*, *180*, 6-10.

Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *Review of Financial Studies*, 33(5), 2223-2273.

Gu, S., Kelly, B., & Xiu, D. (2021). Autoencoder asset pricing models. *Journal of Econometrics*, *222*(1), 429-450.

Han, S. (2022). Is liquidity risk priced in cryptocurrency markets? *Applied Economics Letters*, in press.

Hanauer, M.X., & Kalsbach, T. 2022. Machine learning and the cross-section of emerging market stock returns. TUM working paper.

Hastie, T., Tibshirani, R., & Friedman, J. (2008). *The Elements of Statistical Learning, second edition*. New York: Springer.

Hoang, L. T., & Baur, D. G. (2022). Loaded for bear: Bitcoin private wallets, exchange reserves and prices. *Journal of Banking & Finance*, 144, 106622.

Hoerl, A. E., & Kennard, R. W. (1970). Ridge regression: Biased estimation for nonorthogonal problems. *Technometrics*, 12(1), 55-67.

Hong, H., Lim, T., & Stein, J. C. (2000). Bad news travels slowly: Size, analyst coverage, and the profitability of momentum strategies. *Journal of Finance*, 55(1), 265-295.

Hou, K., Xue, C., & Zhang, L. (2020). Replicating anomalies. *Review of Financial Studies*, 33(5), 2019-2133.

Israel, R., & Moskowitz, T. J. (2013). The role of shorting, firm size, and time on market anomalies. *Journal of Financial Economics*, 108(2), 275-301.

Jacobs, H. (2016). Market maturity and mispricing. *Journal of Financial Economics*, 122(2), 270-287.

Jacobs, H., & Müller, S. (2020). Anomalies across the globe: Once public, no longer existent? *Journal of Financial Economics*, 135(1), 213-230.

Jia, B., Goodell, J. W., & Shen, D. (2022a). Momentum or reversal: Which is the appropriate third factor for cryptocurrencies? *Finance Research Letters*, *45*, 102139.

Jia, Y., Liu, Y., & Yan, S. (2021). Higher moments, extreme returns, and cross–section of cryptocurrency returns. *Finance Research Letters*, *39*, 101536.

Jia, Y., Simkins, B. J., Xu, Z., & Zhang, R. (2022). Psychological anchoring effect, the cross section of cryptocurrency returns, and cryptocurrency market anomalies. Available at SSRN 4170936.

Jones, C. S., & Mo, H. (2021). Out-of-sample performance of mutual fund predictors. *Review of Financial Studies*, 34(1), 149-193.

Kaya, O., & Mostowfi, M. (2022). Low-volatility strategies for highly liquid cryptocurrencies. *Finance Research Letters*, *46*, 102422.

Kelly, B. T., Malamud, S., Zhou, K. (2021). The virtue of complexity in machine learning portfolios. Swiss Finance Institute Research Paper No. 21-90. Available at SSRN: https://ssrn.com/abstract=3984925 or http://dx.doi.org/10.2139/ssrn.3984925.

Kelly, B. T., Pruitt, S., & Su, Y. (2019). Characteristics are covariances: A unified model of risk and return. *Journal of Financial Economics*, *134*(3), 501-524.

Kelly, B., & Pruitt, S. (2013). Market expectations in the cross-section of present values. *Journal of Finance*, 68(5), 1721-1756.

Kelly, B., & Pruitt, S. (2015). The three-pass regression filter: A new approach to forecasting using many predictors. *Journal of Econometrics*, 186(2), 294-316.

Keloharju, M., Linnainmaa, J. T., & Nyberg, P. (2016). Return seasonalities. *Journal of Finance*, *71*(4), 1557-1590.

Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.

Lam, F. E. C., & Wei, K. J. (2011). Limits-to-arbitrage, investment frictions, and the asset growth anomaly. *Journal of Financial Economics*, 102(1), 127-149.

Leippold, M., Wang, Q., & Zhou, W. (2022). Machine learning in the Chinese stock market. *Journal of Financial Economics*, *145*(2A), 64-82.

Leirvik, T. (2021). Cryptocurrency returns and the volatility of liquidity. *Finance Research Letters*, 44, 102031.

Lewellen, J., & Nagel, S. (2006). The conditional CAPM does not explain asset-pricing anomalies. *Journal of Financial Economics*, 82(2), 289-314.

Li, Y., Urquhart, A., Wang, P., & Zhang, W. (2021). MAX momentum in cryptocurrency markets. *International Review of Financial Analysis*, *77*, 101829.

Liebi, L. J. (2022). Is there a value premium in cryptoasset markets? *Economic Modelling*, *109*, 105777.

Lin, C. H., Yen, K. C., & Cheng, H. P. (2021). Lottery-like momentum in the cryptocurrency market. *The North American Journal of Economics and Finance*, *58*, 101552.

Linnainmaa, J. T., & Roberts, M. R. (2018). The history of the cross-section of stock returns. *Review of Financial Studies*, 31(7), 2606-2649.

Liu, W., Liang, X., & Cui, G. (2020). Common risk factors in the returns on cryptocurrencies. *Economic Modelling*, *86*, 299-305.

Liu, Y., & Tsyvinski, A. (2021). Risks and returns of cryptocurrency. *Review of Financial Studies*, *34*(6), 2689-2727.

Liu, Y., Tsyvinski, A., & Wu, X. (2022). Common risk factors in cryptocurrency. *Journal of Finance*, 77(2), 1133-1177.

Liu, Y., Tsyvinski, A., & Wu, Xi. (2021). Accounting for cryptocurrency value. Available at SSRN 3951514.

Llorente, G., Michaely, R., Saar, G., & Wang, J. (2002). Dynamic volume-return relation of individual stocks. *Review of Financial Studies*, 15(4), 1005-1047.

Long, H., Zaremba, A., Demir, E., Szczygielski, J. J., & Vasenin, M. (2020). Seasonality in the cross-section of cryptocurrency returns. *Finance Research Letters*, *35*, 101566.

Lucchini, L., Alessandretti, L., Lepri, B., Gallo, A., & Baronchelli, A. (2020). From code to market: Network of developers and correlated returns of cryptocurrencies. *Science advances*, *6*(51), eabd2204.

McLean, R. D. (2010). Idiosyncratic risk, long-term reversal, and momentum. *Journal of Financial and Quantitative Analysis*, 45(4), 883-906.

McLean, R. D., & Pontiff, J. (2016). Does academic research destroy stock return predictability?. *Journal of Finance*, 71(1), 5-32.

Miller, M. H., & Scholes, M. S. (1982). Dividends and taxes: Some empirical evidence. *Journal of Political Economy*, *90*(6), 1118-1141.

Mohrschladt, H. (2021). The ordering of historical returns and the cross-section of subsequent returns. *Journal of Banking & Finance*, *125*, 106064.

Newey, W.K., & West, K.D. (1987). A simple positive definite, heteroscedasticity, and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703-708.

Novy-Marx, R., & Velikov, M. (2016). A taxonomy of anomalies and their trading costs. *Review of Financial Studies*, 29(1), 104-147.

Ozdamar, M., Akdeniz, L., & Sensoy, A. (2021). Lottery-like preferences and the MAX effect in the cryptocurrency market. *Financial Innovation*, *7*(1), 1-27.7

Pagnotta, E., & Buraschi, A. (2018). An equilibrium valuation of bitcoin and decentralized network assets. Available at SSRN 3142022.

Petropoulos, F., Apiletti, D., Assimakopoulos, V., Babai, M. Z., Barrow, D. K., Taieb, S. B., ... & Ziel, F. (2022). Forecasting: theory and practice. *International Journal of Forecasting*, *38*(3), 705-871.

Pontiff, J. (1996). Costly arbitrage: Evidence from closed-end funds. *The Quarterly Journal of Economics*, 111(4), 1135-1151.

Pontiff, J. (2006). Costly arbitrage and the myth of idiosyncratic risk. *Journal of Accounting and Economics*, 42(1-2), 35-52.

Rad, H., Low, R. K. Y., Miffre, J., & Faff, R. W. (2021). The commodity risk premium and neural networks. Available at SSRN 3816170.

Sadka, R., & Scherbina, A. (2007). Analyst disagreement, mispricing, and liquidity. *Journal of Finance*, 62(5), 2367-2403.

Schwert, G. W. (2003). Anomalies and market efficiency. Handbook of the Economics of Finance, 1, 939-974.

Shen, D., Urquhart, A., & Wang, P. (2020). A three-factor pricing model for cryptocurrencies. *Finance Research Letters*, *34*, 101248.

Stambaugh, R. F., Yu, J., & Yuan, Y. (2012). The short of it: Investor sentiment and anomalies. *Journal of Financial Economics*, 104(2), 288-302.

Stambaugh, R. F., Yu, J., & Yuan, Y. (2015). Arbitrage asymmetry and the idiosyncratic volatility puzzle. *Journal of Finance*, 70(5), 1903-1948.

Struck, C., & Cheng, E. (2020). The cross section of commodity returns: A nonparametric approach. *Journal of Financial Data Science*, *2*(3), 86-103.

Timmermann, A. (2006). Forecast combinations. *Handbook of Economic Forecasting*, *1*, 135-196.

Tobek, O., & Hronec, M. (2021). Does it pay to follow anomalies research? Machine learning approach with international evidence. *Journal of Financial Markets*, *56*, 100588.

Tzouvanas, P., Kizys, R., & Tsend-Ayush, B. (2020). Momentum trading in cryptocurrencies: Short-term returns and diversification benefits. *Economics Letters*, *191*, 108728.

Yang, D., & Zhang, Q. (2000). Drift-independent volatility estimation based on high, low, open, and close prices. *Journal of Business*, 73(3), 477-492.

Zaremba, A., Bilgin, M. H., Long, H., Mercik, A., & Szczygielski, J. J. (2021). Up or down? Short-term reversal, momentum, and liquidity effects in cryptocurrency markets. *International Review of Financial Analysis*, *78*, 101908.

Zaremba, A., Umutlu, M., & Maydybura, A. (2020). Where have the profits gone? Market efficiency and the disappearing equity anomalies in country and industry returns. *Journal of Banking & Finance*, 121, 105966.

Zhang, W., & Li, Y. (2020). Is idiosyncratic volatility priced in cryptocurrency markets?. *Research in International Business and Finance*, *54*, 101252.

Zhang, W., & Li, Y. (2021). Liquidity risk and expected cryptocurrency returns. *International Journal of Finance & Economics.*, in press.

Zhang, W., Li, Y., Xiong, X., & Wang, P. (2021). Downside risk and the cross-section of cryptocurrency returns. *Journal of Banking & Finance*, 133, 106246.

Zou, H., & Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology), 67*(2), 301-320.

Figure 1. A Snapshot of the Research Sample

The figure provides an overview of the sample of cryptocurrencies used in this study. Panel A presents the time-series evolution of the number of cryptocurrencies included in the sample. Panel B illustrates the cross-sectional distribution of the number of daily observations available per cryptocurrency. The red vertical line indicates the average number of observations across all assets in the sample. Panel C displays the sample's aggregate market capitalization (in USD billion) over time. The sample comprises 573 unique cryptocurrencies; the study period runs from July 1, 2017 to July 6, 2022.

Panel A: Time-series composition | Panel B: Cross-sectional composition | Panel C: Market capitalization

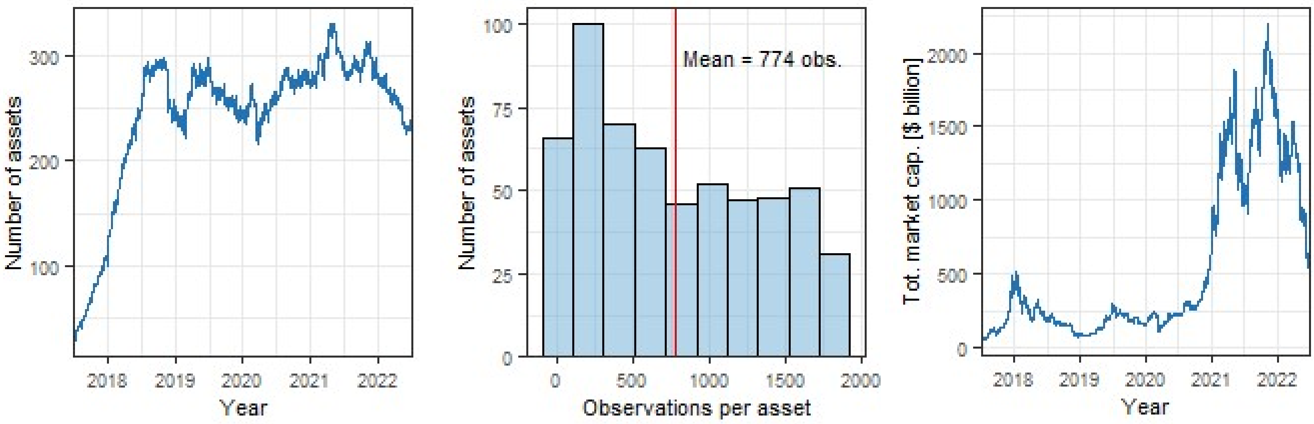

Figure 2. Variable Importance per Characteristic

The figure presents the rankings of 34 return predictors covered in the study in terms of their average total model contribution. The variable importance (VI) is computed as the reduction of the overall OOS R^2 resulting from excluding a given variable from the model. We consider 11 machine learning models: ordinary least squares (OLS), partial least squares (PLS), the least absolute shrinkage and selection operator (LASSO), elastic net (ENET), support vector machine (SVM), gradient-boosted regression trees (GBRT), random forests (RF), feed-forward neural networks with one and two hidden layers (NN1, NN2), and forecast combination (COMB). The color gradients represent the VI rank; the dark blue (white) represents the most influential (least influential) characteristics. The sample comprises 574 cryptocurrencies; the total study period is from July 1, 2017 to July 6, 2022; and the testing period starts on 17 January 2018.

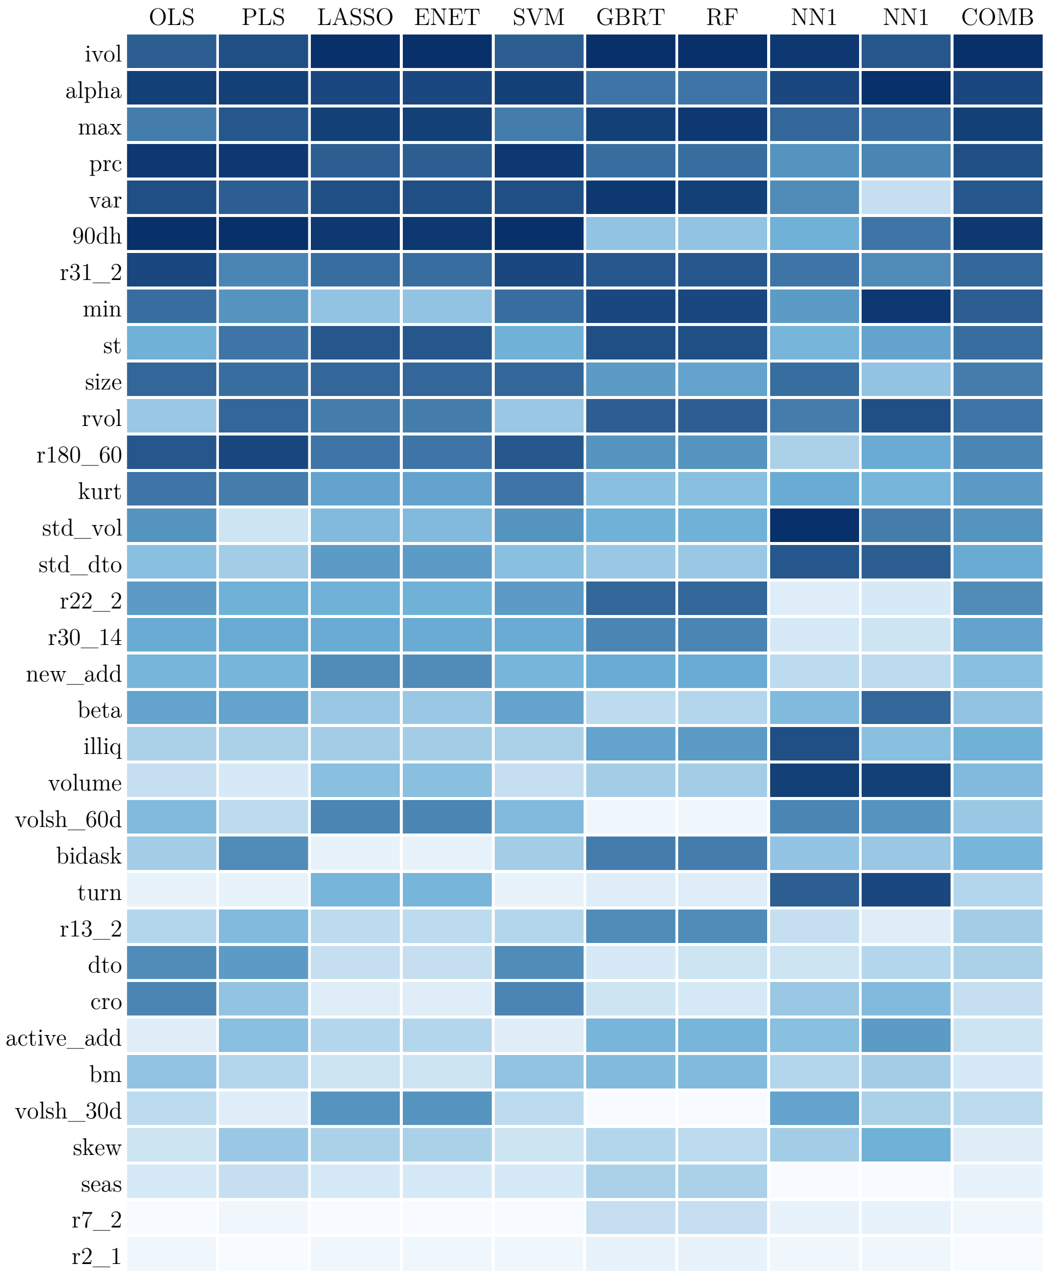

Figure 3. Variable Importance in Different Models

The figure presents the importance of the top 10 variables in the 10 machine learning models considered in this study: ordinary least squares (OLS), partial least squares (PLS), the least absolute shrinkage and selection operator (LASSO), elastic net (ENET), support vector machine (SVM), gradient boosted regression trees (GBRT), random forests (RF), feed-forward neural networks with one and two hidden layers (NN1, NN2, 3), and forecast combination (COMB). The variable importance (VI) is computed as the reduction of the overall OOS R^2 resulting from excluding a given variable from the model. The panels display the reduction in R^2 from setting all values of a given variable to zero in the training sample. VI is averaged across all the training samples and rescaled to sum to 1. The sample comprises 574 cryptocurrencies; the total study period is from July 1, 2017 to July 6, 2022; and the testing period starts on 17 January 2018.

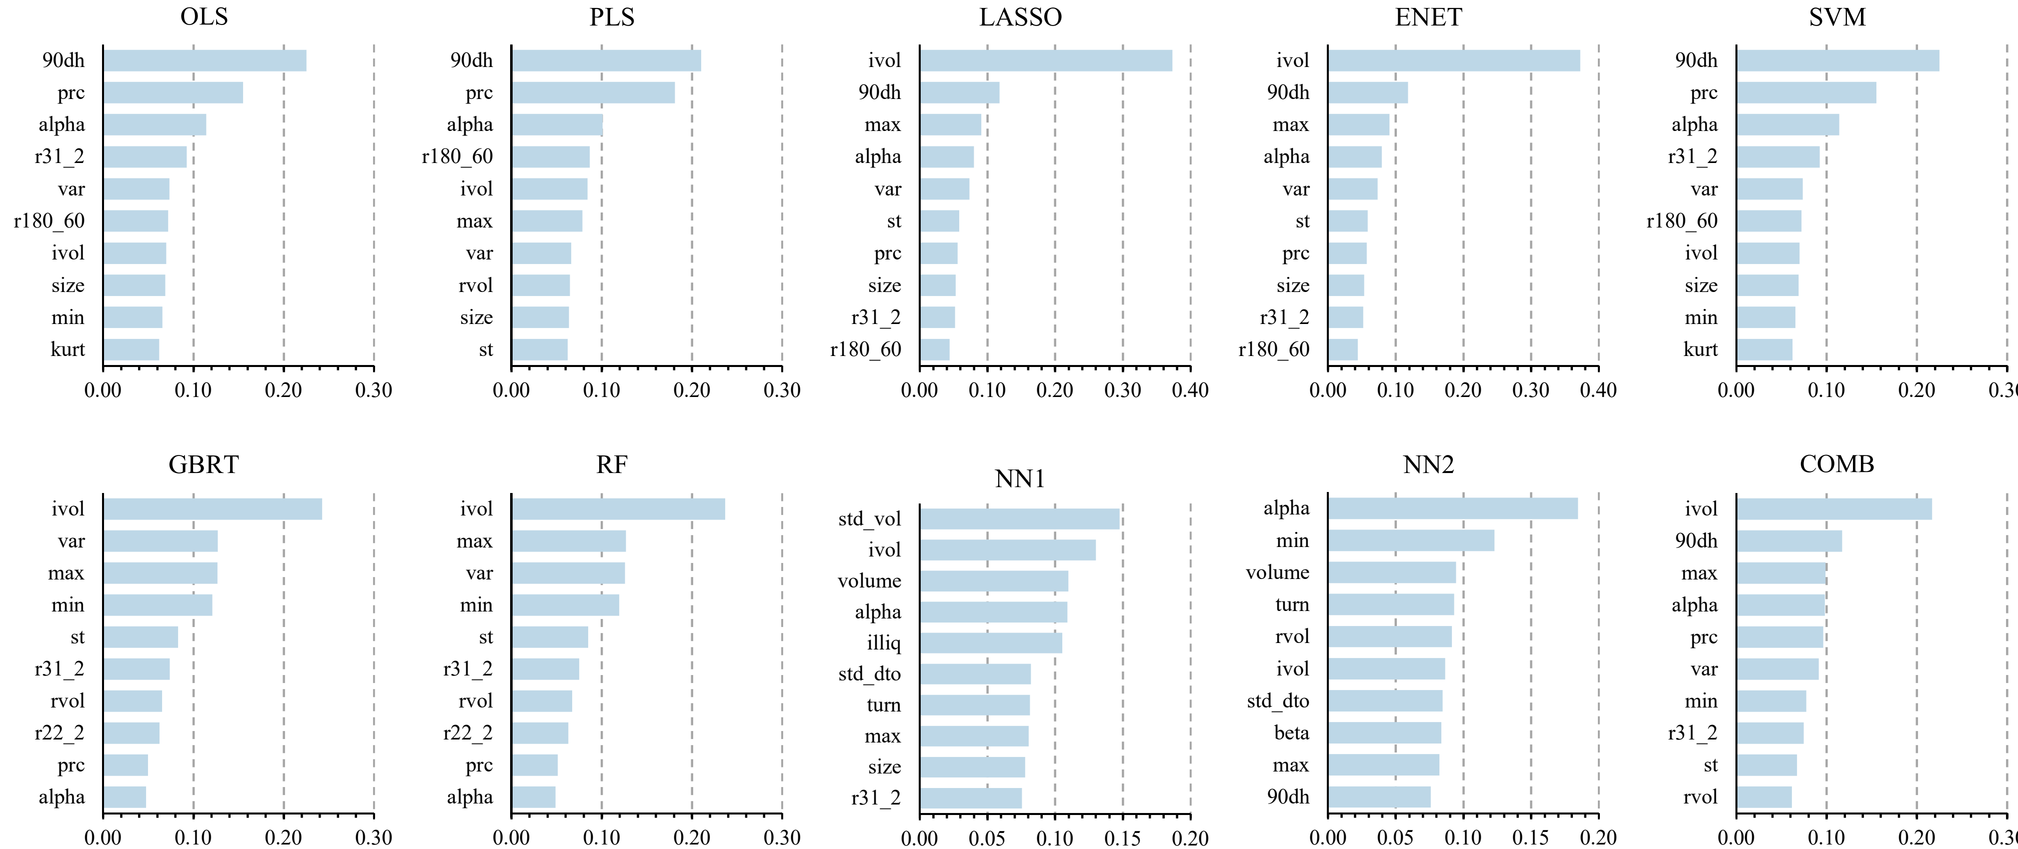

Figure 4. Aggregate Importance of Top Predictors

The figure illustrates the aggregate contribution of the top three, five, and 10 predictors for the machine learning models used in this study: ordinary least squares (OLS), partial least squares (PLS), the least absolute shrinkage, and selection operator (LASSO), elastic net (ENET), support vector machine (SVM), gradient boosted regression trees (GBRT), random forests (RF), feed-forward neural networks with one and two hidden layers (NN1, NN2, 3), and forecast combination (COMB). The variable importance (VI) is computed as the reduction of the overall OOS R^2 resulting from excluding a given variable from the model. VI is averaged across all the training samples and rescaled to 1. The aggregate importance is obtained as in Bali et al. (2022). The sample comprises 574 cryptocurrencies; the total study period is from July 1, 2017 to July 6, 2022; and the testing period starts on 17 January November 2018.

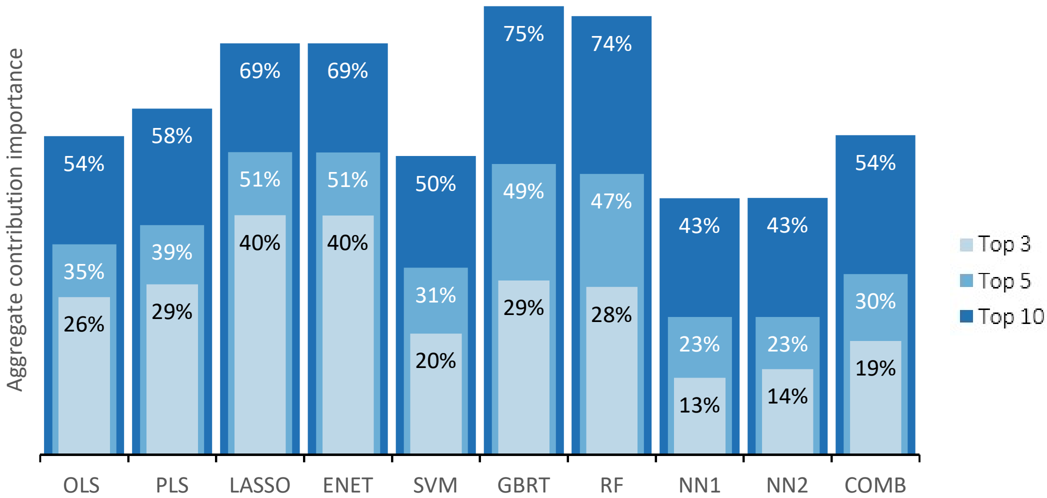

Figure 5. Variable Importance per Category

The figure illustrates the aggregate importance of six categories of cryptocurrency characteristics, as classified in Table 1, in terms of their overall model contribution. We consider 10 machine learning models: ordinary least squares (OLS), partial least squares (PLS), the least absolute shrinkage and selection operator (LASSO), elastic net (ENET), support vector machine (SVM), gradient-boosted regression trees (GBRT), random forests (RF), feed-forward neural networks with one and two hidden layers (NN1, NN2, 3), and forecast combination (COMB). The variable importance (VI) is computed as the reduction of the overall OOS R^2 resulting from excluding a given variable from the model. VI is averaged across all the training samples and is rescaled to sum to 1. The VIs per category are aggregated following the method in Bali et al. (2022). The sample comprises 574 cryptocurrencies; the total study period is from July 1, 2017 to July 6, 2022; and the testing period starts on 17 January November 2018.

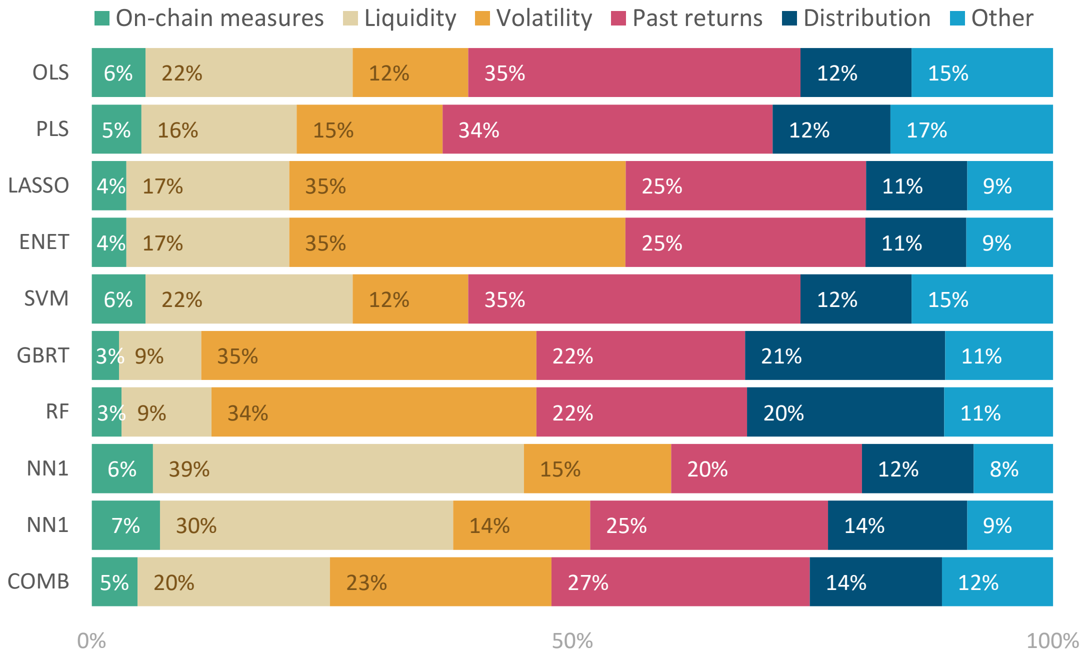

Figure 6. Abnormal Returns on Quintile Machine Learning Portfolios

The figure presents the weekly alphas from the three-factor model of Liu et al. (2022) on quintile portfolios formed on machine learning forecasts. We sort the cryptocurrencies into quintiles based on predictions from 10 models: ordinary least squares (OLS), partial least squares (PLS), the least absolute shrinkage and selection operator (LASSO), elastic net (ENET), support vector machine (SVM), gradient boosted regression trees (GBRT), random forests (RF), feed-forward neural networks with one or two hidden layers (NN1, NN2), and forecast combination (COMB). High (Low) is the cryptocurrency quintile with the highest (lowest) return forecast. The blue bars indicate the portfolios with the lowest expected returns. All alphas are reported in percentage terms. Panels A and B report the results for equal- and value-weighted portfolios, respectively. The sample comprises 574 cryptocurrencies; the total study period is from July 1, 2017 to July 6, 2022; and the testing period starts on 17 January November 2018.

Panel A: Equal-weighted strategies

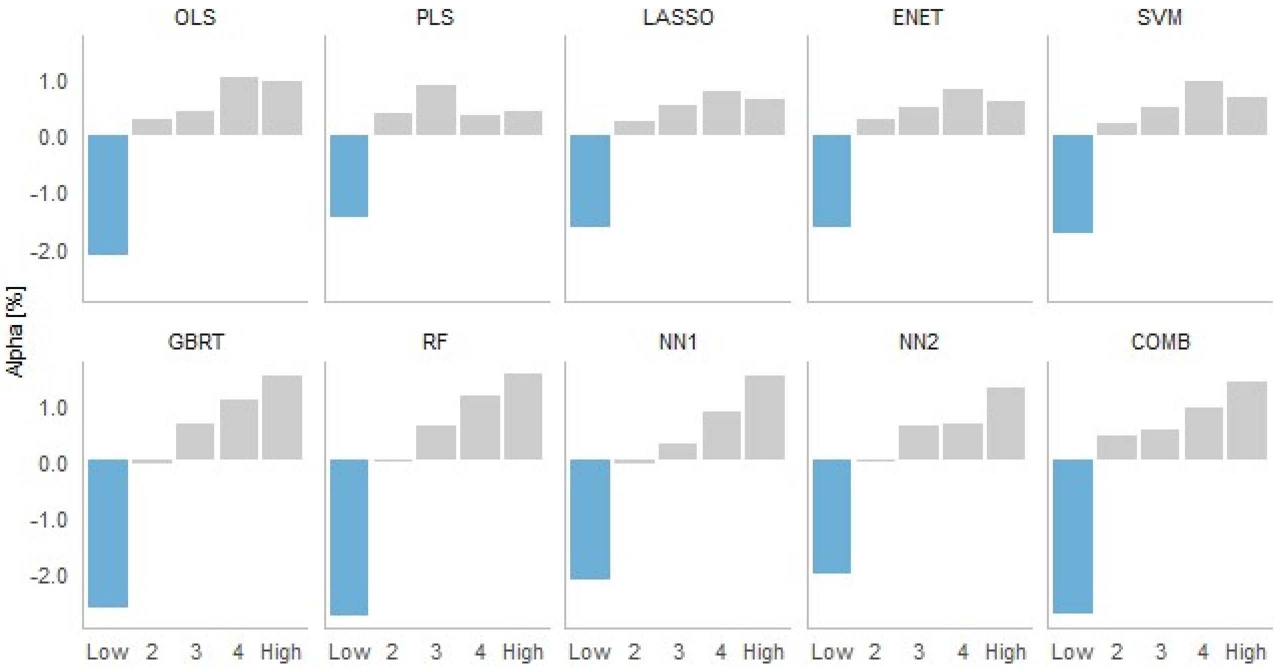

Panel B: Value-weighted strategies

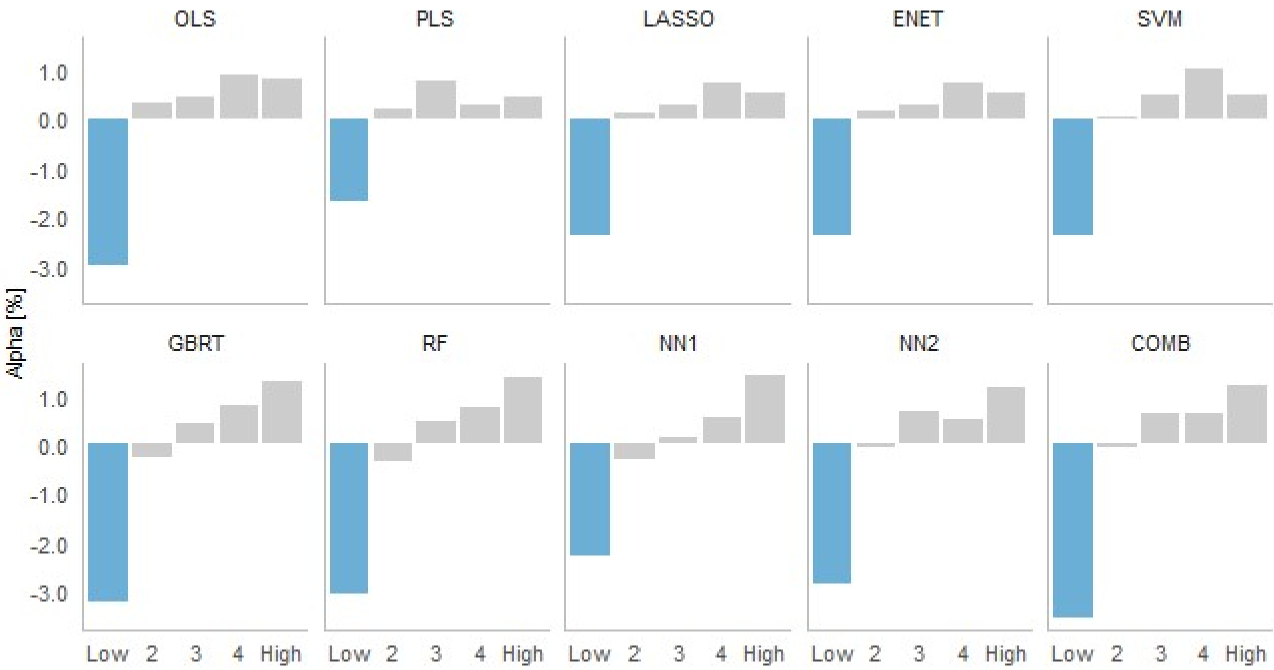

Figure 7. Cumulative Returns on Long-Short Machine Learning Portfolios

The figure presents the cumulative returns on long-short machine learning portfolios. We sort the cryptocurrencies into quintiles based on predictions from 10 models: ordinary least squares (OLS), partial least squares (PLS), the least absolute shrinkage and selection operator (LASSO), elastic net (ENET), support vector machine (SVM), gradient boosted regression trees (GBRT), random forests (RF), feed-forward neural networks with one or two hidden layers (NN1, NN2), and forecast combination (COMB). The reported strategies assume a long (short) position in the quintile of cryptocurrencies with the highest (lowest) expected return. The strategies are rebalanced weekly. Panels A and B report the results for equal- and value-weighted portfolios, respectively. The sample comprises 574 cryptocurrencies; the total study period is from July 1, 2017 to July 6, 2022; and the testing period starts on 17 January November 2018. The blue line represents the strategy indicated in the headline, and the grey lines display other strategies for comparison.

Panel A: Equal-weighted portfolios

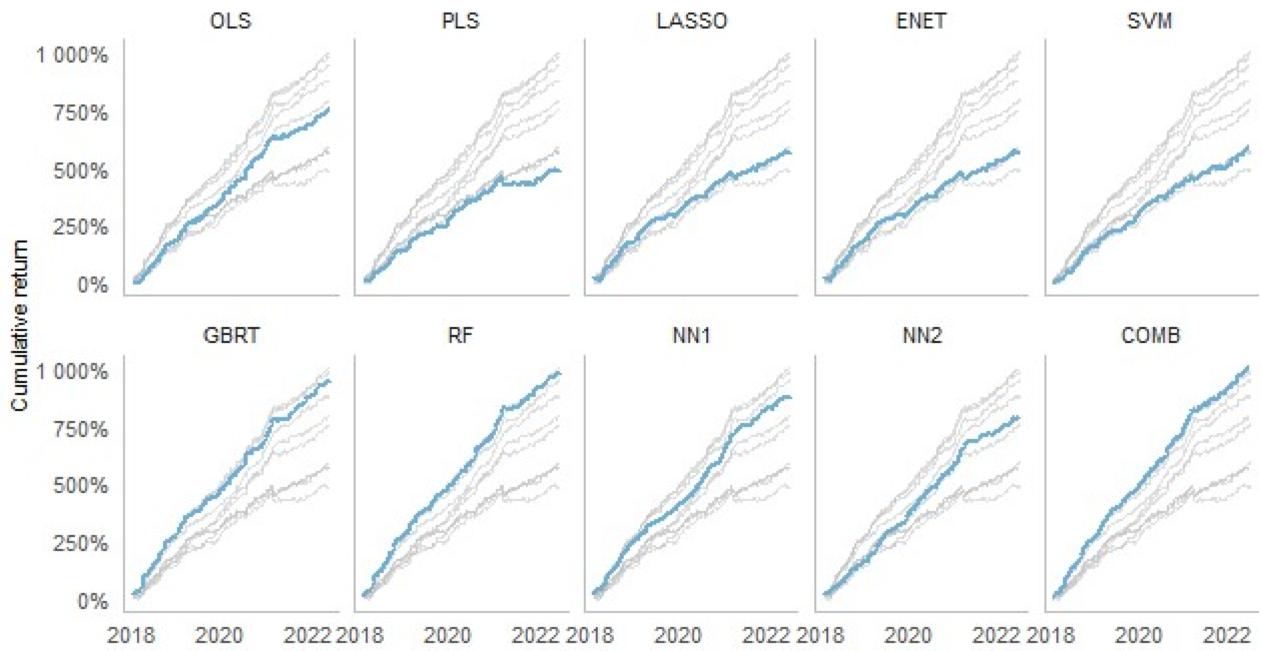

Panel B: Value-weighted portfolios

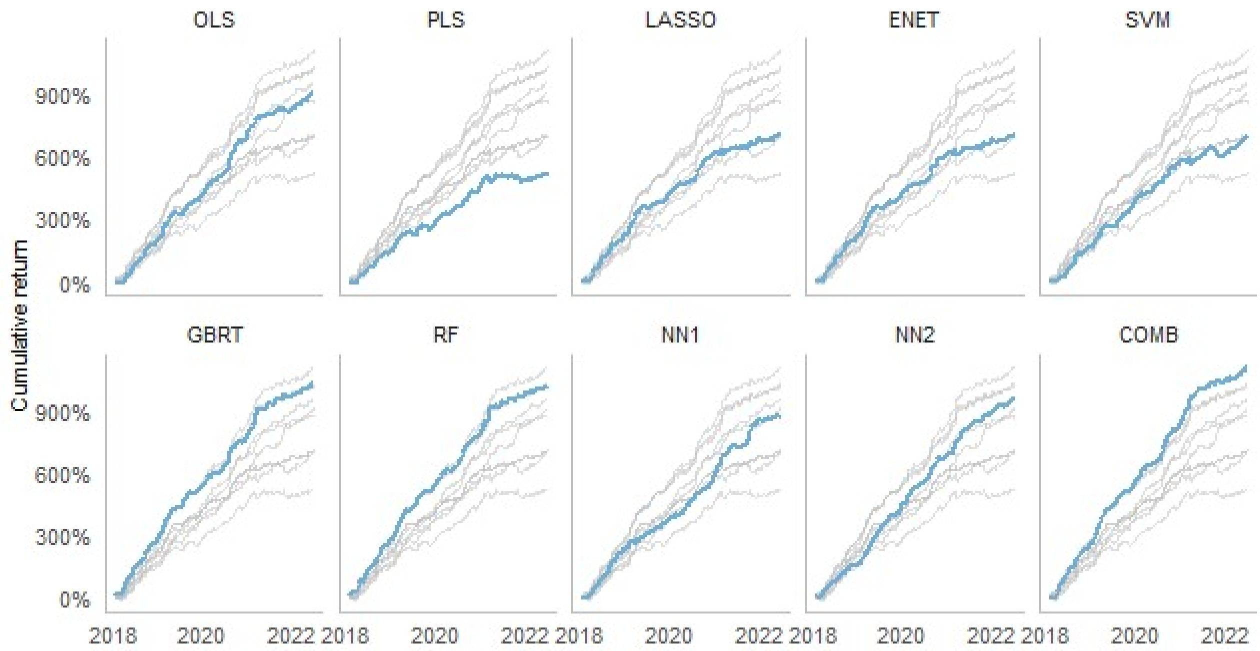

Table 1. Cryptocurrency Characteristics

The table presents the 34 cryptocurrency characteristics used as inputs to machine learning models. *No.* is the running number, and *Symbol* indicates the acronym used to denote the variable in the paper. The table spans two pages.

No.	Characteristic	Symbol	Definition
			On-chain measures
(1)	New addresses	new_add	The number of unique addresses appearing for the first time in a transaction of the native coin in the network (Liu et al., 2021).
(2)	Active addresses	active_add	The number of unique addresses that were active in the network, either as a sender or receiver (Pagnotta & Buraschi, 2018). The calculations are limited to the addresses that were active in successful transactions.
(3)	Network-to-market value	bm	As in Pagnotta and Burashi (2018), the network-to-market value is calculated as the cumulative number of unique addresses divided by the total market value (see Market value for the calculation details).
			Liquidity
(4)	Trading volume	volume	The total dollar value of all native tokens transferred across wallets - both across and within centralized exchanges.
(5)	Market value	size	The available supply times the market price in USD (Liu et al., 2022). The currently available supply is the current supply minus the coins that have been burned.
(6)	Bid-ask spread	bidask	A bid-spread estimation calculated based on 30 days of OHLC data as an average of two approximations by Corwin and Schultz (2012) and Abdi and Ranaldo (2017).
(7)	Illiquidity ratio	illiq	The price impact measure of Amihud (2002), which is calculated as the average 90-day ratio of the absolute value of daily returns over the daily trading volume measured in USD.
(8)	Turnover	turn	The last day's dollar trading volume (see *Trading volume*) over the market value (see *Market Value*) (Datar et al., 2018).
(9)	Detrended turnover	dto	Similar to Garfinkel (2009), the detrended turnover is calculated in two steps. First, we compute the daily excess turnover as the daily turnover (see *Turnover)* minus the value-weighted average daily market turnover. Second, we detrend the obtained value by its 180-day median.
(10)	Turnover volatility	std_dto	Residuals' standard deviation from a regression of daily turnover on a constant using a 30-day estimation period.
(11)	Trading volume volatility	std_vol	Residuals' standard deviation from a regression of daily trading volume on a constant that uses a 30-day estimation period.
(12)	Volume shock (30 days)	volsh_30d	Log-deviation of trading volume from its rolling 30-day average, as in Llorente et al. (2002) and Babiak et al. (2022).
(13)	Volume shock (60 days)	volsh_60d	Log-deviation of trading volume from its rolling 60-day average, as in Llorente et al. (2002) and Babiak et al. (2022).
			Volatility
(14)	Realized volatility	rvol	Daily realized volatility calculated based on 30 days of OHLC prices using the estimator of Yang and Zhang (2000).
(15)	CAPM beta	beta	The market beta from the Capital Asset Pricing Model was using a trailing 30-day period. As in Levellen and Nagel (2006), the beta is calculated as the sum of two slope coefficients from the regression of daily cryptocurrency returns on the contemporaneous and one-day-lagged market excess returns. The market portfolio return is the value-weighted average return of all cryptocurrencies in the sample.
(16)	Idiosyncratic risk	ivol	The standard deviation of the residuals from the regression of daily excess cryptocurrency returns on the daily market portfolio excess returns estimated using a trailing

			30-day period. The market portfolio return is the value-weighted average return of all cryptocurrencies in the sample.
(17)	Value-at-risk	var	The historical empirical value-at-risk computed as the 5th percentile of daily returns over a rolling 90-day period.
			Past returns
(18)	Daily reversal	r2_1	Return on the previous day.
(19)	7-day momentum	r7_2	Cumulative return from seven to two days before return prediction.
(20)	13-day momentum	r13_2	Cumulative return from 13 to two days before return prediction.
(21)	22-day momentum	r22_2	Cumulative return from 22 to two days before return prediction.
(22)	31-day momentum	r31_2	Cumulative return from 31 to two days before return prediction.
(23)	Intermediate moment.	r30_14	Cumulative return from 30 to 14 days before return prediction.
(24)	Long-term reversal	r180_60	Cumulative return from 180 to 60 days before return prediction.
(25)	Closeness to the 90-day high	90dh	Following the logic of George and Hwang (2004), the closeness to the 90-day high is the last day's price over the maximum price over the previous 90 days. The estimation period follows Babiak and Bianchi (2022).
(26)	CAPM alpha	alpha	An intercept from the regression of daily excess cryptocurrency returns on the daily market portfolio excess returns, which is estimated using a trailing 30-day period. The market portfolio return is the value-weighted average return of all cryptocurrencies in the sample.
			Distribution
(27)	Skewness	skew	The skewness of the daily return distribution calculated over a rolling 90-day period.
(28)	Kurtosis	kurt	The kurtosis of the daily return distribution calculated over a rolling 90-day period.
(29)	Maximum daily return	max	The maximum daily return over the last 30 days.
(30)	Minimum daily return	min	The minimum daily return over the last 30 days.
			Other
(31)	Salience theory	st	The salience theory variable is calculated closely following the multistep procedure in Cosemans and Frehen (2021), which uses a rolling 30-day estimation period. We use the market portfolio return as the reference rate and set the parameters θ=0.1 and δ=0.7.
(32)	Chronological return ordering	cro	As in Mohrschladt (2021), the chronological return ordering variable is calculated as the correlation between daily returns over the last 30 days and the corresponding number of trading days until the end of the rolling 30-day estimation window.
(33)	Seasonality	seas	The average same-weekday return calculated over a rolling 20-week period.
(34)	Price	prc	Cryptocurrency price at the end of the previous day.

Table 2. Statistical Properties of Cryptocurrency Characteristics

The table presents the basic statistical properties of the 34 cryptocurrency characteristics used as inputs to machine learning models. The explanation of variable symbols seen in the leftmost column is available in Table 1. The reported values are calculated using a pooled sample of all daily observations. The sample comprises 573 unique cryptocurrencies, and the study period runs from July 1, 2017 to July 6, 2022.

	Mean	Standard deviation	Skewness	Percentiles						
				1st	5th	25th	50th	75th	95th	99th
On-chain measures										
new_add	2907.93	28672.65	13.78	0.00	0.00	1.00	7.00	29.00	516.00	63510.05
active_add	6998.01	64080.47	12.34	0.00	0.00	8.00	29.00	107.00	1286.00	153414.20
bm	0.01	0.11	629.18	0.00	0.00	0.00	0.00	0.01	0.04	0.09
Liquidity										
volume ($mln)	11.97	115.25	25.34	0.00	0.00	0.01	0.06	0.59	22.12	239.50
size ($mln)	2263.34	33343.77	23.61	1.06	1.43	4.55	15.76	69.24	1095.29	18721.31
bidask (%)	4.61	14.47	21.72	0.00	0.00	0.57	2.30	4.97	14.41	34.25
illiq	1.29	171.49	150.64	0.00	0.00	0.00	0.00	0.00	0.01	0.27
turn (%)	2.58	7.07	6.11	0.00	0.00	0.05	0.36	1.82	12.29	37.13
dto (%)	0.49	5.22	5.59	-9.63	-2.70	-0.35	0.00	0.32	5.01	21.44
std_dto (%)	1.69	3.56	3.78	0.00	0.01	0.08	0.34	1.40	8.83	18.64
std_vol ($mln)	6.45	52.43	18.72	0.00	0.00	0.01	0.06	0.51	15.27	130.20
Volatility										
rvol (%)	15.71	39.33	16.10	2.70	4.36	7.03	9.98	14.73	32.60	102.02
beta	0.97	0.82	-0.13	-1.22	-0.11	0.61	0.99	1.34	2.02	3.03
ivol (%)	8.43	6.27	1.98	1.10	2.32	4.25	6.54	10.51	21.21	31.59
volsh_30d	-0.59	1.56	-2.21	-6.78	-3.47	-1.02	-0.31	0.21	1.21	2.19
volsh_60d	-0.71	1.66	-1.97	-7.08	-3.80	-1.24	-0.42	0.20	1.26	2.26
var (%)	-13.85	8.72	-3.73	-49.01	-28.07	-15.68	-11.90	-9.27	-6.13	-3.26
Past returns										
r2_1 (%)	0.26	10.94	1.33	-26.95	-14.10	-4.39	-0.21	3.89	16.21	38.71
r7_2 (%)	0.32	21.82	2.37	-49.16	-27.69	-10.05	-1.10	8.07	33.22	73.88
r13_2 (%)	0.79	33.92	4.17	-64.50	-39.27	-16.03	-2.61	11.66	51.80	115.55
r22_2 (%)	1.81	49.20	5.48	-78.92	-50.51	-22.86	-4.99	15.16	75.54	176.93
r31_2 (%)	3.45	66.11	8.22	-86.94	-58.03	-28.64	-7.03	18.46	96.45	239.36
r30_14 (%)	1.43	42.57	5.11	-73.29	-45.67	-20.03	-3.81	13.86	65.35	149.95
r180_60 (%)	30.14	275.17	22.08	-99.76	-85.57	-55.12	-23.26	33.42	285.07	844.23
90dh (%)	59.10	65.96	521.00	11.87	21.47	39.66	57.85	77.52	98.24	112.78
alpha (%)	0.19	2.03	-1.56	-4.66	-1.92	-0.62	0.02	0.83	3.23	6.29
Distribution										
skew	0.60	1.14	-0.10	-2.64	-1.02	0.01	0.52	1.14	2.55	3.84
kurt	4.26	5.70	3.07	-0.83	-0.11	0.93	2.30	5.30	15.41	27.71
max (%)	25.26	19.26	1.59	2.87	6.41	11.97	19.08	31.90	68.50	92.94
min (%)	-19.48	12.99	-2.26	-72.61	-44.67	-23.24	-16.07	-11.46	-6.77	-3.44
Other										
st	0.00	0.05	-7.90	-0.14	-0.05	-0.02	0.00	0.01	0.04	0.08
cro	0.00	0.17	-0.03	-0.41	-0.28	-0.11	0.00	0.11	0.28	0.40
seas (%)	-0.73	0.33	-1.93	-1.96	-1.31	-0.86	-0.67	-0.52	-0.34	-0.22
prc ($)	159.70	2241.00	22.84	0.00	0.00	0.01	0.06	0.51	32.33	1904.61

Table 3. Prediction Performance of Machine Learning Models

The table presents the prediction performance of the 10 machine learning models used in this study: ordinary least squares (OLS), partial least squares (PLS), the least absolute shrinkage and selection operator (LASSO), elastic net (ENET), support vector machine (SVM), gradient boosted regression trees (GBRT), random forests (RF), feed-forward neural networks with one or two hidden layers (NN1, NN2), and forecast combination (COMB). The reported measures are out-of-sample R^2 coefficients (R^2_{OOS}), calculated as by Gu et al. (2020), as well as the average weekly Pearson ($\bar{\rho}_P$) and Spearman ($\bar{\rho}_S$) correlation coefficients. R^2_{OOS} is expressed in percentage terms. The sample comprises 574 cryptocurrencies; the total study period is from July 1, 2017 to July 6, 2022; and the testing period starts on 17 January 2018.

	OLS	PLS	LASSO	ENET	SVM	GBRT	RF	NN1	NN2	COMB
R^2_{OOS}	0.243	-0.097	-2.234	-2.234	0.067	0.787	0.779	0.826	0.772	0.742
$\bar{\rho}_P$	0.074	0.046	0.054	0.054	0.054	0.106	0.107	0.067	0.069	0.096
$\bar{\rho}_S$	0.072	0.052	0.063	0.063	0.085	0.084	0.083	0.062	0.077	0.092

Table 4. Univariate Portfolio Sorts on Machine Learning Model Predictions

The table presents the average weekly returns on machine learning strategies for cryptocurrency markets. We sort the cryptocurrencies into quintiles based on predictions from 10 models: ordinary least squares (OLS), partial least squares (PLS), the least absolute shrinkage and selection operator (LASSO), elastic net (ENET), support vector machine (SVM), gradient boosted regression trees (GBRT), random forests (RF), feed-forward neural networks with one or two hidden layers (NN1, NN2), and forecast combination (COMB). High (Low) is the cryptocurrency quintile with the highest (lowest) return forecast, and High-Low is the long-short strategy that buys (sells) the top (bottom) quintile. α_1, α_3, and α_6 denote alphas from the one-, three-, and six-factor asset pricing models—respectively. All returns and alphas are reported in percentage terms. The numbers in parentheses are Newey and West's (1987) t-statistics. Panels A and B report the results for equal- and value-weighted portfolios, respectively. The sample comprises 574 cryptocurrencies; the total study period is from July 1, 2017 to July 6, 2022; and the testing period starts on 17 January November 2018.

	OLS	PLS	LASSO	ENET	SVM	GBRT	RF	NN1	NN2	COMB
	Panel A: Equal-weighted portfolios									
Low	-2.71	-2.09	-2.26	-2.27	-2.40	-3.16	-3.30	-2.74	-2.66	-3.37
2	-0.52	-0.40	-0.52	-0.49	-0.52	-0.74	-0.65	-0.73	-0.73	-0.30
3	-0.30	0.11	-0.17	-0.19	-0.09	-0.04	-0.08	-0.43	-0.10	-0.22
4	0.38	-0.22	0.17	0.18	0.28	0.42	0.49	0.27	0.14	0.35
High	0.55	0.01	0.19	0.17	0.14	0.92	0.94	1.03	0.75	0.94
High-Low	3.26	2.10	2.46	2.45	2.54	4.08	4.23	3.77	3.42	4.32
	(10.03)	*(5.78)*	*(7.40)*	*(7.28)*	*(8.70)*	*(10.45)*	*(10.90)*	*(11.46)*	*(11.63)*	*(12.74)*
α_1	3.19	2.04	2.39	2.38	2.49	4.00	4.15	3.72	3.40	4.23
	(10.65)	*(6.47)*	*(7.40)*	*(7.36)*	*(7.97)*	*(11.75)*	*(12.12)*	*(12.19)*	*(11.51)*	*(14.19)*
α_3	3.07	1.86	2.27	2.26	2.44	4.14	4.28	3.63	3.32	4.12
	(9.99)	*(5.74)*	*(6.77)*	*(6.73)*	*(7.55)*	*(11.96)*	*(12.39)*	*(11.57)*	*(10.82)*	*(13.35)*
α_6	2.92	1.76	2.07	2.05	2.27	3.94	4.08	3.66	3.15	3.96
	(8.79)	*(4.96)*	*(5.74)*	*(5.67)*	*(6.43)*	*(10.39)*	*(10.66)*	*(10.78)*	*(9.45)*	*(11.88)*
	Panel B: Value-weighted portfolios									
Low	-3.32	-2.05	-2.77	-2.77	-2.82	-3.51	-3.39	-2.70	-3.27	-3.84
2	-0.22	-0.42	-0.33	-0.28	-0.63	-0.72	-0.75	-0.79	-0.72	-0.57
3	-0.15	0.12	-0.34	-0.37	0.02	-0.04	-0.01	-0.42	0.13	-0.12
4	0.36	-0.18	0.18	0.20	0.55	0.28	0.26	0.16	0.15	0.29
High	0.57	0.16	0.24	0.23	0.14	0.88	0.95	1.00	0.81	0.88
High-Low	3.89	2.21	3.02	3.01	2.96	4.40	4.34	3.70	4.08	4.72
	(8.25)	*(4.98)*	*(6.71)*	*(6.66)*	*(6.48)*	*(8.59)*	*(8.75)*	*(8.24)*	*(10.46)*	*(9.06)*
α_1	3.83	2.17	2.97	2.96	2.92	4.33	4.28	3.62	4.08	4.64
	(10.64)	*(5.70)*	*(7.43)*	*(7.38)*	*(6.70)*	*(10.07)*	*(10.05)*	*(9.80)*	*(10.93)*	*(11.43)*
α_3	3.79	2.11	2.91	2.90	2.88	4.50	4.43	3.70	4.04	4.73
	(10.26)	*(5.33)*	*(7.00)*	*(6.96)*	*(6.40)*	*(10.33)*	*(10.26)*	*(9.65)*	*(10.43)*	*(11.38)*
α_6	3.27	1.60	2.12	2.10	2.23	3.67	3.66	3.75	3.49	3.98
	(8.27)	*(3.76)*	*(4.88)*	*(4.82)*	*(4.62)*	*(8.05)*	*(8.05)*	*(9.00)*	*(8.50)*	*(9.13)*

Table 5. Returns on Machine Learning Portfolios in Subperiods

The table presents the average weekly returns on machine learning strategies for cryptocurrency markets. We sort the cryptocurrencies into quintiles based on predictions from 10 models: ordinary least squares (OLS), partial least squares (PLS), the least absolute shrinkage and selection operator (LASSO), elastic net (ENET), support vector machine (SVM), gradient boosted regression trees (GBRT), random forests (RF), feed-forward neural networks with one or two hidden layers (NN1, NN2), and forecast combination (COMB). High (Low) is the cryptocurrency quintile with the highest (lowest) return forecast, and *High-Low* is the long-short strategy that buys (sells) the top (bottom) quintile. α_3 denotes alphas from the three-factor asset pricing models. All returns and alphas are reported in percentage terms. The numbers in parentheses are Newey and West's (1987) *t*-statistics. The sample comprises 574 cryptocurrencies; the total study period is from July 1, 2017 to July 6, 2022; and the testing period starts on 17 January 2018. Panels A and B report the results for the first and second half of the full testing period; i.e., January 17, 2018 to April 7, 2020 and April 8, 2020 to July 6, 2022.

	OLS	PLS	LASSO	ENET	SVM	GBRT	RF	NN1	NN2	COMB
	Panel A: First half (17 January 2018 to 7 April 2020)									
Low	-5.34	-4.16	-5.26	-5.27	-5.19	-5.75	-5.81	-4.07	-5.42	-6.27
2	-1.89	-2.12	-1.93	-1.89	-2.21	-2.38	-2.20	-2.40	-2.23	-1.86
3	-1.73	-1.38	-1.67	-1.72	-1.68	-1.88	-1.92	-2.27	-1.19	-1.79
4	-1.16	-2.13	-1.73	-1.69	-1.11	-1.19	-1.17	-1.53	-1.75	-1.12
High	-1.15	-1.31	-1.22	-1.23	-1.46	-0.54	-0.46	-0.61	-0.88	-0.78
High-Low	4.19	2.85	4.03	4.04	3.74	5.21	5.35	3.46	4.55	5.49
	(7.01)	*(4.65)*	*(6.40)*	*(6.42)*	*(6.63)*	*(8.46)*	*(9.45)*	*(7.76)*	*(9.56)*	*(8.65)*
α_3	4.22	2.87	4.06	4.06	3.75	5.23	5.38	3.47	4.56	5.52
	(8.98)	*(5.77)*	*(7.46)*	*(7.43)*	*(7.12)*	*(10.42)*	*(10.59)*	*(7.12)*	*(10.14)*	*(10.93)*
	Panel B: Second half (8 April 2020 to 6 July 2022)									
Low	-1.32	0.05	-0.31	-0.29	-0.46	-1.29	-1.00	-1.35	-1.14	-1.43
2	1.44	1.28	1.26	1.31	0.95	0.92	0.68	0.81	0.77	0.71
3	1.42	1.60	0.97	0.97	1.71	1.78	1.87	1.41	1.45	1.54
4	1.87	1.75	2.06	2.06	2.20	1.74	1.68	1.83	2.05	1.68
High	2.27	1.62	1.70	1.69	1.72	2.30	2.34	2.60	2.48	2.53
High-Low	3.59	1.57	2.01	1.98	2.18	3.59	3.34	3.95	3.63	3.96
	(4.96)	*(2.58)*	*(3.64)*	*(3.57)*	*(3.21)*	*(4.56)*	*(4.37)*	*(5.08)*	*(6.01)*	*(4.99)*
α_3	3.47	1.47	1.90	1.88	2.00	3.41	3.17	3.60	3.68	3.75
	(6.30)	*(2.54)*	*(3.30)*	*(3.25)*	*(2.90)*	*(4.89)*	*(4.69)*	*(6.59)*	*(6.18)*	*(5.92)*

Table 6. Machine Learning Performance Over Different Forecast Horizons

The table presents the performance of machine learning models for different forecast horizons: one day (Panels A.1 and B.1) and 14 days (Panels A.2 and B.2). We consider 10 machine learning models: ordinary least squares (OLS), partial least squares (PLS), the least absolute shrinkage and selection operator (LASSO), elastic net (ENET), support vector machine (SVM), gradient boosted regression trees (GBRT), random forests (RF), feed-forward neural networks with one or two hidden layers (NN1, NN2), and forecast combination (COMB). Panel A reports measures of prediction performance: out-of-sample R^2 coefficients (R^2_{OOS}), calculated as in Gu et al. (2020), as well as the average weekly Pearson ($\bar{\rho}_P$) and Spearman ($\bar{\rho}_S$) correlation coefficients. R^2_{OOS} is expressed in percentage terms. Panel B reports the weekly returns on long-short portfolios formed on prediction from the machine learning models. The strategies buy (sell) a quintile of cryptocurrencies with the highest (lowest) predicted return. High-Low R is the mean return, and α_3 denotes alphas from the three-factor asset pricing models. The portfolios are value-weighted, and the security holding period is consistent with the forecast horizon (i.e., one or 14 days). All returns and alphas are reported in percentage terms. The numbers in parentheses are Newey and West's (1987) *t*-statistics. The sample comprises 574 cryptocurrencies; the total study period is from July 1, 2017 to July 6, 2022; and the testing period starts on 17 January 2018.

Panel A: Prediction accuracy

	OLS	PLS	LASSO	ENET	SVM	GBRT	RF	NN1	NN2	COMB
	Panel A.1: One-day forecast horizon									
R^2_{OOS}	-0.094	0.002	-0.219	-0.219	-0.207	-0.038	-0.050	-0.018	-0.032	0.002
$\bar{\rho}_P$	0.017	0.014	0.012	0.012	0.006	0.023	0.023	0.018	0.016	0.022
$\bar{\rho}_S$	0.014	0.013	0.006	0.006	0.044	0.019	0.019	0.013	0.007	0.023
	Panel A.2 14-day forecast horizon									
R^2_{OOS}	-0.125	-0.632	-5.278	-5.279	0.133	0.679	0.665	0.954	1.022	1.453
$\bar{\rho}_P$	0.100	0.063	0.072	0.072	0.088	0.129	0.132	0.131	0.114	0.139
$\bar{\rho}_S$	0.095	0.066	0.078	0.078	0.118	0.103	0.105	0.109	0.105	0.125

Panel B: Long-Short Portfolio Performance

	OLS	PLS	LASSO	ENET	SVM	GBRT	RF	NN1	NN2	COMB
	Panel B1: One-day forecast horizon									
High-Low R	2.38	1.55	2.00	2.00	1.34	2.71	2.84	1.99	1.57	2.94
	(5.42)	*(3.37)*	*(3.72)*	*(3.72)*	*(2.57)*	*(5.72)*	*(6.08)*	*(4.31)*	*(3.62)*	*(6.44)*
High-Low α_3	2.29	1.60	2.00	2.00	1.48	2.58	2.73	1.85	1.40	2.87
	(4.84)	*(3.41)*	*(3.64)*	*(3.65)*	*(2.79)*	*(4.91)*	*(5.21)*	*(3.67)*	*(2.69)*	*(5.96)*
	Panel B.2: 14-day forecast horizon									
High-Low R	3.84	2.28	3.06	3.04	3.20	3.96	4.17	4.27	4.30	4.82
	(7.39)	*(4.53)*	*(6.86)*	*(6.83)*	*(6.05)*	*(6.85)*	*(7.62)*	*(7.58)*	*(10.51)*	*(7.78)*
High-Low α_3	3.86	2.19	3.08	3.06	3.02	4.18	4.35	4.18	4.44	4.89
	(10.04)	*(4.33)*	*(7.37)*	*(7.32)*	*(6.04)*	*(9.12)*	*(10.21)*	*(9.07)*	*(11.49)*	*(10.90)*

Table 7. Interactions Between Machine Learning Predictions and Limits to Arbitrage

The table reports the average slope coefficients from weekly predictive cross-sectional regressions following Fama and MacBeth (1973); where the dependent variable is a weekly cryptocurrency return, and the independent variables include expected returns (*E(R)*) from the forecast combination (COMB) model, proxies of limits to arbitrage, and their interaction terms. The limits to arbitrage are represented by dummy variables that take a value of one if a given proxy is higher than a weekly median, or zero otherwise. The leftmost column indicates the dummies for four different proxies for limits to arbitrage: idiosyncratic volatility ($ivol_D$), illiquidity ($illiq_D$), bid-ask spread ($bidask_D$), and a composite measure (lim_D). The numbers in parentheses are Newey and West's (1987) adjusted *t*-statistics. The bottom row presents the average weekly cross-sectional adjusted R^2 coefficient ($\bar{R}^2$) expressed in percentage terms. The sample comprises 574 cryptocurrencies; the total study period is from July 1, 2017 to July 6, 2022; and the testing period starts on 17 January 2018.

	(1)	(2)	(3)	(4)	(5)
E(R)	1.352	0.701	0.699	0.692	0.872
	(12.61)	*(5.31)*	*(5.24)*	*(5.29)*	*(6.47)*
$ivol_D$		0.014			
		(2.17)			
E(R)×$ivol_D$		1.167			
		(5.96)			
$bidask_D$			0.014		
			(2.09)		
E(R)×$bidask_D$			1.165		
			(5.93)		
$illiq_D$				0.014	
				(2.11)	
E(R)×$illiq_D$				1.177	
				(6.11)	
lim_D					0.007
					(1.32)
E(R)×lim_D					0.846
					(4.24)
$\bar{R}^2$	1.38	1.88	1.88	1.86	1.82

Table 8. Bivariate Portfolio Sorts on Limits to Arbitrage and Predicted Returns

The table presents the portfolios from bivariate sorts on different measures of limits to arbitrage and expected returns (*E(R)*) from the forecast combination (COMB) model. In the first step, we sort the cryptocurrencies on their limits to arbitrage using four different proxies: idiosyncratic volatility (*ivol*), illiquidity (*illiq*), bid-ask spread (*bidask*), and a composite measure (*lim*). Subsequently—within each of these subsets—we sort cryptocurrencies into Low *E(R)*, Medium *E(R)*, and High *E(R)* terciles (as indicated in the top row) based on the COMB predictions. All portfolios are rebalanced weekly. Panel A displays the average weekly returns on the considered portfolios; it additionally presents the average returns (*H-L R*) and three-factor model (Liu et al., 2022) alphas (*High-Low* α_3) on the spread portfolio that buys (sells) the coins with the High (Low) *E(R)*. The left and right sections concern equal- and value-weighted strategies, respectively. The bottom rows of each panel present the differences in returns on the *High-Low* portfolios between the *High* and *Low* limits to arbitrage terciles. All returns and alphas are reported in percentages. The numbers in parentheses are Newey and West's (1987) adjusted *t*-statistics. Panel B pertains to the relative importance of stocks in different portfolios. Its left side reports the average market value of coins in different categories (in $ billion), and the right side illustrates the average proportion (in percentage) of the total market capitalization allocated therein. The sample comprises 574 cryptocurrencies; the total study period is from July 1, 2017 to July 6, 2022; and the testing period starts on 17 January 2018.

Panel A: Average returns

	Equal-weighted portfolios							Value-weighted portfolios						
	Low E(R)	Medium E(R)	High E(R)	High-Low R	t-stat$_R$	High-Low α_3	t-stat$_\alpha$	Low E(R)	Medium E(R)	High E(R)	High-Low R	t-stat$_R$	High-Low α_3	t-stat$_\alpha$
Panel A.1: Idiosyncratic volatility														
Low ivol	-0.77	-0.11	0.51	1.27	*(3.76)*	0.83	*(2.54)*	-0.98	-0.06	0.64	1.62	*(3.71)*	1.17	*(3.01)*
Medium ivol	-1.46	-0.30	0.72	2.19	*(5.58)*	2.13	*(5.46)*	-1.86	-0.40	0.58	2.44	*(4.07)*	2.68	*(5.37)*
High ivol	-4.16	0.10	0.81	4.97	*(9.59)*	4.84	*(9.99)*	-5.82	0.12	0.65	6.47	*(9.41)*	6.56	*(9.50)*
High-Low ivol				3.69	*(5.51)*	4.01	*(6.92)*				4.85	*(5.55)*	5.38	*(6.69)*
Panel A.2: Illiquidity														
Low illiq	-1.50	-0.44	0.88	2.39	*(6.76)*	2.00	*(6.23)*	-1.63	-0.24	0.82	2.45	*(6.03)*	2.21	*(5.84)*
Medium illiq	-1.86	-0.06	0.77	2.63	*(6.67)*	2.43	*(6.04)*	-1.92	0.12	0.50	2.43	*(4.82)*	2.35	*(4.88)*
High illiq	-3.02	0.18	0.38	3.40	*(7.30)*	3.35	*(7.63)*	-5.65	-0.11	0.53	6.19	*(6.98)*	6.30	*(7.65)*
High-Low illiq				1.02	*(1.65)*	1.35	*(2.40)*				3.74	*(3.74)*	4.10	*(4.37)*
Panel A.3: Bid-ask spread														
Low bidask	-1.70	0.01	0.84	2.54	*(7.79)*	2.31	*(6.36)*	-2.21	-0.19	0.99	3.20	*(8.20)*	3.08	*(7.48)*
Medium bidask	-1.57	-0.29	0.64	2.21	*(6.07)*	1.97	*(5.74)*	-1.47	-0.03	0.54	2.01	*(3.82)*	1.94	*(4.13)*
High bidask	-3.12	-0.04	0.56	3.68	*(7.79)*	3.51	*(7.81)*	-4.61	-0.10	0.36	4.97	*(5.50)*	4.96	*(6.43)*
High-Low bidask				1.13	*(1.83)*	1.20	*(2.00)*				1.77	*(1.83)*	1.88	*(2.08)*
Panel A.4: Composite measure														
Low lim	-1.07	-0.12	0.68	1.75	*(5.16)*	1.37	*(4.27)*	-1.30	-0.13	0.66	1.96	*(5.15)*	1.67	*(4.56)*
Medium lim	-1.97	-0.18	1.02	2.99	*(8.26)*	2.93	*(7.85)*	-2.50	0.05	0.85	3.35	*(5.08)*	3.42	*(5.62)*
High lim	-3.36	-0.03	0.32	3.67	*(6.64)*	3.47	*(7.16)*	-5.29	-0.37	0.24	5.53	*(6.06)*	5.54	*(6.82)*
High-Low lim				1.92	*(2.88)*	2.10	*(3.55)*				3.57	*(3.76)*	3.87	*(4.33)*

Panel B: Average cryptocurrency capitalization and market proportions

	Average market cap ($ bil)			Average market proportion (%)			
	Low E(R)	Medium E(R)	High E(R)	Low E(R)	Medium E(R)	High E(R)	Sum
Panel B.1: Idiosyncratic volatility							
Low ivol	0.35	2.17	15.85	1.79	11.11	80.92	93.82
Mid ivol	0.08	0.22	0.40	0.38	1.10	2.02	3.50
High ivol	0.05	0.14	0.34	0.25	0.70	1.73	2.68
Panel B.2: Illiquidity							
Low illiq	0.43	2.44	16.38	2.18	12.47	83.65	98.30
Mid illiq	0.03	0.06	0.15	0.16	0.32	0.75	1.23
High illiq	0.02	0.02	0.05	0.09	0.12	0.26	0.47
Panel B.3: Bid-ask spread							
Low bidask	0.37	1.71	8.91	1.87	8.75	45.47	56.09
Mid bidask	0.08	0.74	6.99	0.42	3.79	35.64	39.85
High bidask	0.03	0.08	0.69	0.14	0.42	3.51	4.06
Panel B.4: Composite measure							
Low lim	0.41	2.36	16.03	2.12	12.08	82.25	96.45
Mid lim	0.04	0.13	0.37	0.22	0.64	1.90	2.75
High lim	0.02	0.04	0.10	0.08	0.20	0.51	0.80

Table 9. Univariate Portfolio Sorts – Practical Perspective

The table concerns practical aspects of the machine learning portfolio formed on cryptocurrency predictions. We sort the cryptocurrencies into quintiles based on predictions from 10 models: ordinary least squares (OLS), partial least squares (PLS), the least absolute shrinkage and selection operator (LASSO), elastic net (ENET), support vector machine (SVM), gradient boosted regression trees (GBRT), random forests (RF), feed-forward neural networks with one or two hidden layers (NN1, NN2), and forecast combination (COMB). *High* (*Low*) is the cryptocurrency quintile with the highest (lowest) return forecast, and *High-Low* is the long-short strategy that buys (sells) the top (bottom) quintile. All strategies are value-weighted and rebalanced weekly. Panel A reports different risks and performance metrics for weekly returns on the *High-Low* portfolios: standard deviation, annualized Sharpe ratio, maximum weekly loss, and maximum drawdown over the study period. Panel B displays the average market capitalization (in $ billion) of cryptocurrencies in different quintiles, while Panel C shows their relative market proportion (in %). Finally, Panel D presents the average weekly portfolio turnover (in %)—interpreted as the share of the portfolio being replaced every week. The sample comprises 574 cryptocurrencies; the total study period is from July 1, 2017 to July 6, 2022; and the testing period starts on 17 January November 2018.

	OLS	PLS	LASSO	ENET	SVM	GBRT	RF	NN1	NN2	COMB
	Panel A: Risk profile of the long-short strategies									
Standard dev.(%)	5.58	5.83	6.14	6.15	6.65	6.64	6.57	5.80	5.67	6.34
Sharpe ratio	5.03	2.73	3.54	3.52	3.20	4.78	4.77	4.60	5.19	5.37
Max 1W loss (%)	-12.21	-24.12	-24.98	-24.98	-33.43	-15.43	-11.12	-11.58	-16.40	-10.74
Max DD (%)	-32.23	-44.28	-46.95	-46.95	-58.83	-37.64	-35.29	-33.35	-40.53	-32.28
	Panel B: Average cryptocurrency capitalization ($ billion)									
Low	0.2	0.1	0.1	0.1	0.2	0.4	0.5	0.6	0.3	0.1
2	0.6	0.3	0.1	0.1	0.4	0.7	1.0	1.2	0.7	0.3
3	1.3	0.7	0.3	0.3	0.6	1.6	1.5	2.6	1.5	1.0
4	2.5	2.4	1.0	1.0	2.2	2.9	3.4	3.5	2.8	1.9
High	6.3	7.4	9.3	9.3	7.5	5.3	4.5	3.0	5.6	7.6
	Panel C: Market proportion (%)									
Low	1.7	1.0	1.0	1.0	1.7	2.1	3.3	9.1	2.7	0.8
2	5.5	4.7	1.6	1.6	3.1	10.4	13.9	12.2	9.1	2.6
3	12.9	10.0	5.1	5.1	6.6	18.4	15.3	22.8	14.1	10.8
4	23.4	28.4	14.9	15.2	19.2	24.0	26.7	29.6	23.4	21.3
High	56.6	55.9	77.4	77.0	69.4	45.1	40.7	26.3	50.6	64.5
	Panel D: Portfolio turnover (%)									
Low	47.2	41.8	46.1	46.2	60.1	50.1	50.8	58.0	55.1	46.8
2	64.4	61.8	60.1	60.1	75.6	62.2	63.9	74.5	71.0	64.8
3	67.7	64.0	61.7	61.8	75.4	67.5	67.7	75.9	72.0	67.9
4	63.8	59.5	56.3	56.3	71.4	63.6	64.2	72.6	68.0	63.0
High	39.1	36.1	32.5	32.6	50.7	42.2	42.7	52.2	46.8	37.2
High - Low	86.3	77.9	78.6	78.8	110.8	92.3	93.5	110.2	101.9	83.9

Table 10. Trading Costs of Machine Learning Strategies

The table presents the impact of trading costs on the performance of long-short cryptocurrency machine-learning portfolios. We consider 10 machine learning models: ordinary least squares (OLS), partial least squares (PLS), the least absolute shrinkage and selection operator (LASSO), elastic net (ENET), support vector machine (SVM), gradient-boosted regression trees (GBRT), random forests (RF), feed-forward neural networks with one or two hidden layers (NN1, NN2), and forecast combination (COMB). The strategies buy (sell) a quintile of cryptocurrencies with the highest (lowest) predicted return. All strategies are rebalanced weekly, and Panels A and B concern equal- and value-weighted portfolios—respectively. Gross return is the average weekly return on a long-short portfolio. Break-even TC (trading cost) is the level of the one-way transaction costs at which the average returns decrease to zero. Net returns are estimated using two approaches: a) assuming a flat transaction fee that equals 30 (40) basis points for the long (short) positions (*Net return, flat TC*); and b) as a sum of half of the effective spread estimated using the methods of Corwin and Schultz (2012) and Abdi and Ranaldo (2017), and a flat fee of 10 basis points (*Net return, variable TC*). All the returns are reported in percentages, and the breakeven costs are expressed in basis points. The numbers in parentheses are *t*-statistics. The sample comprises 574 cryptocurrencies; the total study period is from July 1, 2017 to July 6, 2022; and the testing period starts on January 17, 2018.

	OLS	PLS	LASSO	ENET	SVM	GBRT	RF	NN1	NN2	COMB
	Panel A: Equal-weighted portfolios									
Gross return	3.26	2.10	2.46	2.45	2.54	4.08	4.23	3.77	3.42	4.32
	(10.03)	*(5.78)*	*(7.40)*	*(7.28)*	*(8.70)*	*(10.45)*	*(10.90)*	*(11.46)*	*(11.63)*	*(12.74)*
Break-even TC (bps.)	197	142	170	169	123	239	243	177	170	277
Net return (flat TC)	2.68	1.58	1.95	1.93	1.81	3.49	3.62	3.02	2.71	3.77
	(8.24)	*(4.35)*	*(5.86)*	*(5.75)*	*(6.20)*	*(8.92)*	*(9.33)*	*(9.19)*	*(9.23)*	*(11.12)*
Net return (variable TC)	1.04	0.35	0.74	0.72	-0.92	1.80	1.82	0.17	0.13	2.38
	(3.22)	*(0.97)*	*(2.24)*	*(2.15)*	*(-3.15)*	*(4.60)*	*(4.69)*	*(0.53)*	*(0.45)*	*(7.02)*
	Panel B: Value-weighted portfolios									
Gross return (%)	3.89	2.21	3.02	3.01	2.96	4.40	4.34	3.70	4.08	4.72
	(8.25)	*(4.98)*	*(6.71)*	*(6.66)*	*(6.48)*	*(8.59)*	*(8.75)*	*(8.24)*	*(10.46)*	*(9.06)*
Break-even TC (bps.)	226	142	192	191	133	238	232	168	200	281
Net return (flat TC)	3.28	1.66	2.45	2.44	2.17	3.74	3.68	2.93	3.36	4.13
	(6.95)	*(3.74)*	*(5.46)*	*(5.41)*	*(4.76)*	*(7.31)*	*(7.42)*	*(6.51)*	*(8.61)*	*(7.92)*
Net return (variable TC)	2.01	0.70	1.22	1.19	-0.22	2.24	2.17	0.83	1.51	2.89
	(4.26)	*(1.59)*	*(2.72)*	*(2.65)*	*(-0.49)*	*(4.37)*	*(4.37)*	*(1.85)*	*(3.87)*	*(5.54)*