The European Journal of Finance

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/rejf20

The dynamics of returns predictability in cryptocurrency markets

Daniele Bianchi, Massimo Guidolin & Manuela Pedio

To cite this article: Daniele Bianchi, Massimo Guidolin & Manuela Pedio (2022): The dynamics of returns predictability in cryptocurrency markets, The European Journal of Finance, DOI: 10.1080/1351847X.2022.2084343

To link to this article: https://doi.org/10.1080/1351847X.2022.2084343

THE EUROPEAN JOURNAL OF FINANCE
https://doi.org/10.1080/1351847X.2022.2084343

Routledge
Taylor & Francis Group

OPEN ACCESS Check for updates

The dynamics of returns predictability in cryptocurrency markets

Daniele Bianchi [a], Massimo Guidolin [b] and Manuela Pedio [c,d]

[a]School of Economics and Finance, Queen Mary University of London, London, UK; [b]Department of Finance, Universitá Bocconi and Baffi CAREFIN, Milan, Italy; [c]School of Accounting and Finance, University of Bristol, Bristol, UK; [d]Baffi CAREFIN, Milan, Italy

ABSTRACT
In this paper, we take a forecasting perspective and compare the information content of a set of market risk factors, cryptocurrency-specific predictors, and sentiment variables for the returns of cryptocurrencies vs traditional asset classes. To this aim, we rely on a flexible dynamic econometric model that not only features time-varying coefficients, but also allows for the entire forecasting model to change over time to capture the time variation in the exposures of major digital currencies to the predictive variables. Besides, we investigate whether the inclusion of cryptocurrencies in an already diversified portfolio leads to additional economic gains. The main empirical results suggest that cryptocurrencies are not systematically predicted by stock market factors, precious metal commodities or supply factors. On the contrary, they display a time-varying but significant exposure to investors' attention. In addition, also because of a lack of predictability compared to traditional asset classes, cryptocurrencies lead to realized expected utility gains for a power utility investor.

ARTICLE HISTORY
Received 8 August 2021
Accepted 3 May 2022

KEYWORDS
Bitcoin; cryptocurrencies; returns predictability; investments; dynamic model averaging

JEL CODES
G11; G12; G17

1. Introduction

The efficient market hypothesis states that financial markets are efficient with respect to a particular information set when prices aggregate all available information. Market efficiency is then synonymous with the inability of investors to make economic, i.e. risk-adjusted, profits based on a given available information set. As such, the question of returns predictability is of tremendous importance for investors, market participants and academics alike. For investors, the presence of return predictability means making optimal investment decisions which may lead to substantial welfare gains. For academics, return predictability, or the lack thereof, has substantial implications for general equilibrium models that are able to accurately describe the risks and returns in financial markets. While existing literature has extensively studied returns predictability in traditional markets, such as equity, commodity and foreign exchange, little attention has been devoted to the extent of returns predictability within the context of cryptocurrency markets. In this paper we empirically address this issue.

Our approach is that the extent to which returns predictability exists in cryptocurrency markets can be proxied by the out-of-sample performance of state-of-the-art predictive models and the risk-adjusted economic gain generated by these models. In particular, in this paper we are driven by the belief that an appealing way to understand the differences and similarities between cryptocurrencies and other conventional assets is to investigate whether their expected returns behave similarly based on a *forecasting framework*. The very fact that cryptocurrency markets are created and traded in a highly-fragmented, multi-platform structure, which is decentralized and granular, adds to the conjecture that they may be separated from traditional asset classes. Yet, their relative predictability compared to traditional asset classes could have first-order implications for portfolio decision making and therefore the rising interest by investors for cryptocurrencies.

CONTACT Daniele Bianchi d.bianchi@qmul.ac.uk

Supplemental data for this article can be accessed here. https://doi.org/10.1080/1351847X.2022.2084343

Empirically, we focus on the predictability of the returns on four major cryptocurrencies, namely Bitcoin (BTC), Ethereum (ETH), Litecoin (LTC) and Ripple (XRP) (see, e.g. Liu and Tsyvinski 2020), as well as five traditional asset classes (US stocks, worldwide developed country stocks ex-US, US investment grade corporate bonds, spot gold, and a long-short positions equivalent to the trade-weighted US dollar exchange rate). The choice of these four cryptocurrencies is dictated by (1) the length of time time series available, (2) the depth of market liquidity (these being majorly traded cryptocurrencies), and (3) the fact that they have been regularly investigated in related research (see, e.g. Cheah et al. 2022). The sample period varies depending on the asset considered with BTC and traditional asset classes being the longest time series considered from January 2011 to January 2021. We assess the predictive power of a variety of crypto market characteristics, stock market predictors, and sentiment variables both from a statistical and an economic perspective. That is, we do not simply focus on whether cryptocurrencies can be successfully predicted, but we assess whether such predictability (if any) can be dynamically exploited to generate positive economic value to investors that already hold diversified portfolios (including stocks, bonds, commodities, and currencies). This allows us to draw conclusions on whether cryptocurrencies represent a segmented asset class.[1]

Notably, in contrast to much of the existing literature that has investigated the predictability of crypto asset returns (see, e.g. Detzel et al. 2021; Cheah et al. 2022), we consider it to be crucial to employ a flexible econometric approach to capture the relationship between asset returns and predictive variables, especially in the case of cryptocurrencies. In fact, when applied to capture linear predictive relationships, standard recursive regression methods suffer from three problems. First, the sensitivity of returns to a change in a given predictor may not necessarily be time invariant in the data generating process. There is a large literature in macroeconomics and finance which documents structural breaks and other sorts of parameter change in many time series (see, among many others, Stock and Watson 1996 and recursive regression methods are poorly designed to capture such parameter change. Second, the number of potentially relevant predictors can be large and not known a priori, that is, the nature and the number of risk factors that explain the dynamics of cryptocurrency returns is characterized by pervasive uncertain. In light of this fact, an ever expanding literature has turned to Bayesian methods, either by performing Bayesian model averaging (BMA) or by automatizing the model selection process. Yet, even in these cases, computational demands can become formidable when the researcher is facing a number of models that grows with the power of the number of predictors, P, i.e. 2^P. Third, the model relevant to a forecasting application may potentially change over time (see, e.g. Pesaran and Timmermann 2005). This implies that not only there is dynamics in a time-series sense, i.e. time-varying parameters, but also in a cross-sectional sense, i.e. the number of parameters that are significant at a given point in time may change throughout the sample. This issue further complicates an already daunting econometric exercise: if the researcher faces 2^P models and, at each point in time, a different predictive model applies, the number of combinations of models which must be estimated in order to forecast the target variable is $2^{T\times P}$. Even in simple prediction exercises, it can be computationally infeasible to forecast by simply going through such an amount of model combinations.

To address this issue, we follow a strategy originally developed by Raftery, Kárnỳ, and Ettler (2010) and then exploited in macroeconomic forecasting applications by Koop and Korobilis (2012), which is commonly referred to as dynamic model averaging (henceforth, DMA). Their approach can also be used for dynamic model selection (DMS) in which a single (potentially varying) model is used to derive forecasts at each point in time. DMA and DMS seem ideally suited to yield maximum flexibility within the problem of predicting cryptocurrency returns as they allow for the predictive model to undergo structural change, at the same time letting the coefficients in each model evolve over time.[2] Interestingly, both DMA and DMS may allow for both gradual or abrupt changes in the role of a predictor. By construction, standard time-varying parameter regressions – especially when the coefficients are assumed to follow a random walk, as it is typical – fail to allow for such abrupt changes and instead force smooth variation of the slope coefficients over time.

1.1. Findings

One of the explanations commonly invoked for the rising interest by investors for cryptocurrencies is their alleged diversification benefits: crypto asset returns are presumed to display low or negative correlation with the returns of traditional asset classes like bonds and stocks. There is some evidence that their inclusion in traditional portfolios may increase expected returns per unit of risk (see Bouri et al. 2018a, 2018b; Corbet et al. 2018; Bouri et al. 2019; Plakandaras, Bouri, and Gupta 2021). Such evidence is mostly relegated to static predictive frameworks and mean-variance investor preferences. We depart from both these assumptions and report four key results.

First, digital currency returns are on average less predictable when compared to other asset classes, including gold and the log-differences in a trade value-weighted US dollar exchange rate. Both recursive, total mean-squared forecast error (MSFE) and relative out-of-sample (OOS) R-square measures show that, while dynamic modeling strategies deliver relatively low MSFEs and positive (often high) OOS R-squares for traditional asset classes, this is not always the case for cryptocurrencies. Second, crypto asset returns can be forecast according to patterns and with a measurable degree of time variation that differ from most other asset classes, including gold, which has been often indicated as the most closely related asset (see Klein, Thu, and Walther 2018). For instance, consistently with the empirical results in Drobetz, Momtaz, and Schröder (2019) and Li and Yi (2019), an attention index – here measured by the rate of growth of the Google searches concerning each cryptocurrency – proves to be a key prediction variable, which is often not the case with traditional asset classes. Third, in recursive asset allocation experiments, cryptocurrencies generate quite substantial, realized OOS economic value (especially when measured in terms of ex-post certainty equivalent returns) when they are added to otherwise traditional asset menus of cash, corporate bonds, US and international stocks, and long-short exchange rate positions. Moreover, the advantage in terms of realized certainty equivalent returns tend to decrease for higher levels of risk aversion. Finally, the value of cryptocurrencies cannot be reduced only to the fact the investors may have access to long positions in Bitcoin, in the sense that also Ethereum, Litecoin, and Ripple appear to contribute to the OOS realized economic value when they are added to the asset menu in addition to Bitcoin.

1.2. Related literature

This paper contributes to two main strands of research. First, it adds to a growing literature that aims at understanding the investment properties of cryptocurrencies. Yermack (2015) and Dyhrberg (2016a) have investigated the diversification properties of Bitcoin within the context of a diversified portfolio and reached opposite conclusions. In particular, Yermack (2015) argues that Bitcoin is uncorrelated with the majority of fiat currencies and is much more volatile, therefore being of limited usefulness for risk management purposes and diversification (see also Baur, Hong, and Lee 2018; Schilling and Uhlig 2019 for a different conclusion). Chuen et al. (2017) place themselves on a middle ground as they show that, although cryptocurrencies are mildly correlated with traditional asset classes so that they expand the efficient frontier, they fail to lead to a large improvement in the utility of mean variance investors. Bianchi and Babiak (2022) uses a large panel of digital currencies and finds that there is no significant relationship between cryptocurrencies and global proxies of risk for equity. Plakandaras, Bouri, and Gupta (2021) show the OOS economic benefit of adding Bitcoin to a stock-bond portfolio under several different allocation strategies. Dyhrberg (2016b) explore the hedging properties of Bitcoin and find that it can be used as a hedge against stocks in the Financial Times Stock Exchange Index as well as the American dollar. Kajtazi and Moro (2019) explore the effects of adding Bitcoin to an optimal portfolio of US, European and Chinese portfolio assets. They show that Bitcoin improves the returns of the portfolio, mostly from increased returns rather than lower volatility, and that Bitcoin has a role in portfolio diversification.

We contribute to this literature in at least three ways. First, we consider a power utility investor that adds cryptocurrencies to an already well-diversified portfolio (including commodities and currencies). The use of a power utility investor has important implications within the context of cryptocurrency markets. As a matter of fact, the significant departure from normality in the dynamics of realized returns make a typical mean-variance framework almost mechanically misspecified. Second, we consider a full-blown dynamic parametric model which explicitly accommodates for time-varying model uncertainty and volatility. The empirical results shows that the ability of dynamically selecting the most suitable forecasting model at each point in time significantly improves the forecasting as well as the economic performance. Third, we show that, ceteris paribus, the time-series predictability of cryptocurrency returns is lower than the predictability of traditional asset classes. This provides evidence of market segmentation from an ex-ante forecasting perspective. In this respect, we complement and extend the existing literature by showing that the returns dynamics of digital currency markets cannot be indeed reconciled with standard predictors.

Second, we contribute to a large literature on time-varying return predictability. Among many others, examples are Pastor and Stambaugh (2009), Rapach, Strauss, and Zhou (2010a), Binsbergen, Jules, and Koijen (2010), Dangl and Halling (2012), Pettenuzzo, Timmermann, and Valkanov (2014), and Johannes, Korteweg, and Polson (2014). We extend this literature by looking at the predictive power of otherwise standard predictors within the context of cryptocurrency markets. Crucially, our objective is not to overthrow existing results from the traditional returns predictability literature, but rather to draw a direct comparison with the existing research on other asset classes in order to better understand the economics of cryptocurrency markets. In this respect, due to the institutional differences, we view our paper as an OOS test of existing evidence on time-varying predictability developed in more traditional financial markets.

2. Research design

Consider a standard, off-the-shelf time varying parameter (TVP) regression model specification of the form (see, e.g. Dangl and Halling 2012; Bianchi, Guidolin, and Ravazzolo 2017; Guidolin, Hansen, and Pedio 2019 among others):

$$r^j_{t+1} = \theta^j_{0,t} + \sum_{p=1}^{P} (\theta^j_{p,t})' z^j_{p,t} + \epsilon^j_{t+1} = (\boldsymbol{\theta}^j_t)' \mathbf{z}^j_t + \epsilon^j_{t+1} \; t = 1, \ldots, T, \quad j = 1, \ldots, J \tag{1}$$

$$\boldsymbol{\theta}^j_t = \boldsymbol{\theta}^j_{t-1} + \boldsymbol{\eta}_t \; t = 1, \ldots, T, \quad j = 1, \ldots, J \tag{2}$$

where $\mathbf{z}^j_t$ is a jth $(P+1) \times 1$ vector of predictors potentially specific to currency/asset j, always including a unit constant to absorb the intercept coefficient and possibly some aggregate bias, $\boldsymbol{\theta}^j_t$ is a $(P+1) \times 1$ of possibly time-varying (state-dependent) coefficients, ϵ^j_t IID $N(0, h^j_t)$, $\boldsymbol{\eta}_t$ IID $N(\mathbf{0}, \mathbf{L}^j_t)$, and the errors ϵ^j_t and $\boldsymbol{\eta}_t$ are mutually independent at all leads and lags and for all assets. This model is general and can be applied to any type of asset returns as far as data are stationary.

Notably, in standard forecasting applications, the state-space model in Equations (1) and (2) is based on the assumption that the same set of explanatory variables is a-priori relevant at all points in time. While this might seem a fairly innocuous assumption, when P is large, the threat that a rich TVP regression over-fits the data appears to be substantial, which may substantially damage the predictive power of the overall framework. To address this issue, following Koop and Korobilis (2012), we overlay on this model a simple technique to let alternative models to hold at different points in time while their coefficients remain time-varying and then average across them to compute predictions.

Suppose one has a set of Q models which are characterized by different subsets of $\mathbf{z}^j_t$ predictors. Denoting these by $\mathbf{z}^j_t(q)$ for $q = 1, 2, \ldots, Q$, the state-space model in Equations (1) and (2) can be written as (we drop the j index for clarity):

$$r_{t+1} = \boldsymbol{\theta}_t{}'(q) \mathbf{z}_t(q) + \epsilon_{t+1}(q) \; t = 1, \ldots, T, \quad q = 1, \ldots, Q \tag{3}$$

$$\boldsymbol{\theta}_t(q) = \boldsymbol{\theta}_{t-1}(q) + \boldsymbol{\eta}_t(q) \; t = 1, \ldots, T, \quad q = 1, \ldots, Q, \tag{4}$$

where $\epsilon_t(q)$ IID $N(0, h_t(q))$, $\boldsymbol{\eta}_t(q)$ IID $N(\mathbf{0}, \mathbf{L}_t(q))$, and the errors are mutually independent at all leads and lags. Let $M_t \in \{1, 2, \ldots, Q\}$ denote the model at time t. Our framework allows to calculate $Pr(M_t = q \,|\, x^t)$ for $q = 1, 2, \ldots, Q$, that is the probability that a given model q is selected at time t based on the available historical information on returns and predictors $x^t = (r^t, \mathbf{z}^t)$. When forecasting the time $t+1$ returns, we can then take the average of $Pr(M_t = q \,|\, x^t)$ for all $q = 1, 2, \ldots, Q$ at a given time t; this is called Dynamic Model Averaging (DMA). Alternatively, we can simply use the model with the highest $Pr(M_t = q \,|\, x^t)$ and only this specific model for forecasting returns at time $t+1$; a method that is dubbed Dynamic Model Selection (DMS). The latter is particularly problematic when the set of possible models become relatively large. In particular, although a natural approach is to let the model selection to be driven by a latent Markov chain so that how predictors enter/leave the model in real time is simply captured by a transition matrix, Π, with elements $\pi_{ij} = \Pr(M_{t+1} = j \mid M_t = i)$, see e.g. Guidolin (2011), inference in such models is computationally infeasible because Π will be considerably large for interesting choices of the number of predictors, P.[3] To tackle the inherent over-parameterization of DMA, we use an approximation to the Kalman filter originally proposed by Raftery, Kárnỳ, and Ettler (2010) and implemented by Koop and Korobilis (2012) for macroeconomic forecasting.

This method depends on two parameters, λ and ψ, which we shall call forgetting factors and set to be slightly below one, to allow decay over time to occur. Their role is best explained in a standard Kalman filter iteration when model switching is (just for illustrative purposes) ignored. The benefit of using this approximation in the model prediction equation is that we do not require an MCMC algorithm to draw transitions between models nor a simulation algorithm over the space of models. A complete description of the DMA model estimation procedure is provided in Appendix A. DMS proceeds instead by selecting the single model with the highest value of the filtered model probability at each point in time. With reference to the conditional heteroskedastic process $h_t^j(q)$, we follow Raftery, Kárnỳ, and Ettler (2010) and simply proceed to plug in its place a consistent estimate that in line with existing applied work using DMA and DMS methods is based on an Exponentially Weighted Moving Average (EWMA) estimator. Appendix A provides all required details.

One comment is in order. Our dynamic framework implies that forecasting is performed in real time. This means that, for all our variables, we use the value which would have been available to the forecaster at the time the forecast was being made with any hindsight bias ruled out.

2.1. Data and descriptive statistics

The empirical analysis is based on weekly returns of four major cryptocurrencies, namely Bitcoin (BTC), Ethereum (ETH), Litecoin (LTC) and Ripple (XRP). The choice to analyze these four crypto assets is not uncommon in the literature (see, e.g. Atsalakis et al. 2019; Liu and Tsyvinski 2020; Hudson and Urquhart 2021) as their fraction of the total market capitalization oscillates between 60 % and 90 % over our sample period. The reason why we focus on weekly returns is twofold: first, some of the key predictors, such as the google search volume index, are only available at weekly frequency. Second, we aim to maintain a certain degree of consistency with the existing literature (see, e.g. Liu and Tsyvinski 2020).

Cryptocurrency price data series are collected from Coinmetrics.io. For Bitcoin, we use data from 1 January 2011 to 31 December 2021. We start in 2011 because there was not much liquidity and trading in earlier years, see Bianchi, Babiak, and Dickerson (2022). The data series for Litecoin ranges from 21 July 2013 to 31 December 2021; the one for Ripple from 12 April 2015 to 31 December 2021; finally, the price series for Ethereum are the shortest, spanning a 5 June 2016 to 31 December 2021 sample. The differences in the length of the samples reflect the different launching dates of these cryptocurrencies. However, in all cases, the samples are relatively short than what may be available, because these minor cryptocurrencies saw trading and liquidity increase only between 2013 and 2015. All samples end in December 2021, which implies that we have at least 5 years and a half of data (that is, at least 291 observations) for all the cryptocurrencies under analysis. In addition, including 2020 and 2021 in the sample period allows us to capture the impact of the Covid-19 crisis on the global economy and financial markets.

As far as the returns on traditional asset classes are concerned, the prices of the closest to expiry futures on the main precious metals (gold, silver, and platinum) are retrieved from investing.com. The aggregate stock returns on both global and developed stock markets are from CRSP. The data on corporate bonds are the returns on the ICE BofAML US Corporate Index, which tracks the performance of US dollar denominated investment-grade rated corporate debt publicly issued in the US domestic market.[4]

We consider a wide range of possible predictors which are either specific to each cryptocurrency or more generally used for macroeconomic forecasting within the context of traditional asset classes. We obtain the (excess) returns of the market, the small-minus-big (SMB), the high-minus-low (HML), the momentum, the robust-minus-weak (RMW) and the conservative-minus-aggressive (CMA) factors for developed countries from Ken French's website (see Fama and French 2015). In addition, we calculate the time-series momentum for the returns on each cryptocurrency as the 12-week compounded net returns with one week of holding period. Next, to proxy for supply factors, we get back to basics and acknowledge that 'mining', or producing, a digital currency requires two inputs [5]: electricity and computer power. As for electricity, we consider two proxies, i.e. the value-weighted stock returns of U.S.-listed electricity firms and the value-weighted stock returns of the China-listed electricity companies. The reason why we include the China proxies is because electricity supply is location-specific whilst China is considered to have the largest coin mining operation among all countries for a non-trivial fraction of our sample (see, e.g. Li et al. 2019).[6]

In addition to electricity supply, we consider the stock returns of the companies that are major manufacturers of either GPU mining chips (Nvidia Corporation and Advanced Micro Devices, Inc., AMD for short) or ASIC mining chips (Taiwan Semiconductor Manufacturing Company, Limited and Advanced Semiconductor Engineering, Inc., TSMC for short), as proxies of increasing computing power (see, e.g. Liu and Tsyvinski 2020). Our choice to use the returns of specific company stocks to proxy for computing power and mining costs derives from an attempt to resort to forward-looking measures of such factors, therefore reflecting current and expected future supply conditions. Finally, we use the growth rate in the Google search volume index (SVI) downloaded from Google as a proxy for investors' attention (similar to Da, Engelberg, and Gao (2011). Earlier literature (see, e.g. Drobetz, Momtaz, and Schröder 2019; Li and Yi 2019) has shown the importance of investors' attention in explaining cryptocurrency returns. A full list of all the variables with details on computation and sources can be found in Table A.1 of the Online Appendix.

Table 1 provides some descriptive statistics for both the returns and the predictors. While the four cryptocurrencies span (partly by construction and partly by choice) different sample periods, we report summary statistics for the remaining portfolio and asset returns as well as for the non-financial predictors with reference to the longest, 2011–2020 sample, in line with the time series of Bitcoin returns. The first panel reports summary statistics for the digital currency returns. A few stylized facts already well known in the literature emerge. All cryptocurrencies carry very high mean weekly returns, in the order of 132% (for Litecoin) to as much as 224% (for Ripple) in annualized terms. This is due the rapid growth in the value of cryptocurrencies experimented between 2014 and 2017 and again in mid-2020. However such high mean returns are countered by two features. First, all cryptocurrencies are characterized by astonishingly high realized weekly volatilities, ranging between 16% and 37 % (which amounts to a volatility of 115–269 % in annualized terms, using standard formulas). Second, on top of such high volatilities, all cryptocurrencies (but Ethereum, for which the recorded series are however shorter than for the rest of our data) display massive excess kurtosis coefficients, between 8 and 133 (see also Baur, Hong, and Lee 2018; Bianchi and Babiak 2022). This means that, even after discounting their high variability as measured by their second moment, crypto returns are plagued by massive tails, which reflects that hugely negative and positive returns are always possible.

One last feature of crypto returns is interesting: their weekly median returns are systematically much lower (even negative in the case of Ripple) vs. weekly mean returns. Even though an analysis of the empirical distributions of cryptocurrency returns reveals that multiple modes are possible, medians that are systematically lower than means are indicative of positive skewness, which indeed appears to be considerable, between 1.1 (Ethereum) and 9.8 (Ripple). In spite of their large variability, all cryptocurrencies are marked by very attractive *weekly* Sharpe ratios, between 0.12 for Ripple (that however displays an extremely appealing right-skewness) to a stunning 0.22 for Bitcoin. Finally, as shown in much earlier literature (see, e.g. Klein, Thu, and Walther 2018), cryptocurrencies are characterized by negative or positive but modest correlations with the returns on the remaining asset classes (the highest pairwise correlation is 0.17 between Litecoin and value-weighted US stock returns). As concerns the pairwise correlations between the cryptocurrencies themselves, Litecoin displays high correlations with Ripple (0.6) and Ethereum (0.52). In contrast, Bitcoin returns have low pairwise correlations with the returns of the rest of the cryptocurrencies. In fact, those correlations amount to 0.15 for Ethereum and Litecoin, and 0.06 for Ripple.

The empirical facts about the other asset/portfolio returns are generally better known, with a few surprises. On the one hand, the value-weighted market index yield an annualized (excess) mean return of approximately 14.9%, annualized volatility of 16.2%, negative skewness (which makes aggregate stock market returns structurally different from cryptocurrency returns) and some non-negligible excess kurtosis that is however much lower than in the case of crypto (with the sole exception of Ethereum). The Sharpe ratio is between half and

Table 1. Summary statistics.

	Sample	# weeks	Mean	Median	St. Dev.	IQ range	Skew	Ex. Kurt.	SR
					Cryptocurrency Returns				
Bitcoin	Jan. 2011 – Dec 2021	574	3.365	1.644	15.663	11.987	1.857	8.312	0.215
Litecoin	July 2013 – Dec 2021	441	2.554	0.327	21.206	12.858	4.266	33.927	0.120
Ripple	April 2015 – Dec 2021	351	4.309	−0.757	37.267	13.196	9.807	133.067	0.116
Ehereum	June 2016 – Dec 2021	291	3.250	1.892	16.432	16.353	1.069	4.023	0.198
					Other Asset/Portfolio Returns				
1-month T-bill	Jan. 2011 – Dec 2021	574	0.010	0.002	0.015	0.015	1.395	0.410	
Inv. Grade US Corp bond	Jan. 2011 – Dec 2021	574	0.098	0.143	0.788	0.697	−3.211	41.989	0.124
USD Trade-Weighted Exchange Rate	Jan. 2011 – Dec 2021	574	0.043	0.077	0.721	0.877	0.530	2.113	0.060
VW CRSP	Jan. 2011 – Dec 2021	574	0.286	0.396	2.260	2.231	−0.684	7.154	0.127
Developed Market Portfolio	Jan. 2011 – Dec 2021	574	0.210	0.346	2.171	1.978	−0.741	7.323	0.097
SMB portfolio	Jan. 2011 – Dec 2021	574	−0.038	0.000	0.752	0.959	−0.890	6.581	−0.051
HML portfolio	Jan. 2011 – Dec 2021	574	−0.069	−0.090	1.154	1.038	0.515	6.180	−0.060
RMW portfolio	Jan. 2011 – Dec 2021	574	0.085	0.080	0.585	0.697	0.081	1.642	0.146
CMA portfolio	Jan. 2011 – Dec 2021	574	−0.022	−0.035	0.619	0.670	0.536	3.086	−0.035
Momentum	Jan. 2011 – Dec 2021	574	0.101	0.180	1.573	1.471	−1.510	9.179	0.064
Bitcoin Momentum	Jan. 2011 – Dec 2021	574	2.364	1.612	5.773	5.920	1.422	4.178	0.410
Litecoin Momentum	July 2013 – Dec 2021	441	1.121	−0.120	6.004	6.129	1.502	2.724	0.187
Ripple Momentum	April 2015 – Dec 2021	351	1.389	0.037	6.738	4.691	2.737	10.209	0.206
Ethereum Momentum	June 2016 – Dec 2021	291	2.224	1.544	5.595	6.243	0.930	1.271	0.398
Gold	Jan. 2011 – Dec 2021	574	0.065	0.125	2.209	2.415	−0.174	2.071	0.029
Platinum spot	Jan. 2011 – Dec 2021	574	−0.044	−0.150	3.377	3.642	0.029	5.489	−0.013
Silver spot	Jan. 2011 – Dec 2021	574	0.031	−0.035	4.121	4.008	−0.723	7.667	0.008
NVIDIA individual stock	Jan. 2011 – Dec 2021	574	0.943	0.745	5.774	6.482	0.647	2.708	0.163
AMD individual stock	Jan. 2011 – Dec 2021	574	0.838	0.434	7.902	8.763	0.662	3.555	0.106
TSMC individual stock	Jan. 2011 – Dec 2021	574	0.454	0.349	3.529	4.433	−0.046	0.682	0.129
Developed (ex US) VW Mkt	Jan. 2011 – Dec 2021	574	0.093	0.180	2.291	2.226	−1.035	12.897	0.040
					Other predictors				
Google BTC searches	Jan. 2011 – Dec 2021	574	6.661	−1.835	44.061	22.442	4.816	35.301	
Google LTC searches	July 2013 – Dec 2021	441	4.016	−2.381	46.262	25.970	4.993	38.155	0.087
Google XRP searches	April 2015 – Dec 2021	351	2.900	0.000	32.375	11.646	6.494	69.050	0.090
Google ETH searches	June 2016 – Dec 2021	291	4.630	0.000	35.911	28.746	2.471	10.415	0.129
Chinese electricity stocks	Jan. 2011 – Dec 2021	574	0.078	0.215	3.006	2.920	−0.913	8.682	0.026
US electricity stocks	Jan. 2011 – Dec 2021	574	0.167	0.245	2.398	2.334	−0.079	15.325	0.070

Notes: The statistics are based on weekly data, expressed in US dollars. All commodity and foreign stock returns are expressed in US dollars. The interquartile range is defined as the difference between the 75th and the 25th sample percentiles. Excess kurtosis is defined as sample kurtosis minus 3. The Sharpe ratio is computed with reference to the weekly yields of US 30-day T-bills.

two-thirds compared to cryptocurrencies, which is largely known. The surprises come from most of the five Fama-French factor-mimicking long-short portfolio: SMB (representing size), HML (value), and CMA (investment) portfolio returns are all negative, either on average or in median terms. It is well known that the so-called ‘smart beta’ factor portfolios generate positive risk premia. However, such a requirement is to be intended to apply on average, unconditionally and over long periods of time: this explains why on our 11-year data sample, it occurs that many such factors fail to return a positive premium. In any event, the RMW (the quality) and Carhart’s momentum factors do return large premia and strong Sharpe ratios.

While gold, platinum and silver appear to have declined or remained stable on average and therefore to have yielded performances that largely differ from those of the cryptocurrencies, the three individual stock return series (especially NVIDIA and AMD, the producers of GPU mining chips) display properties that tend to be consistent with the general summary statistics for the cryptocurrencies. Finally, the Google searches for the word that correspond to the cryptocurrencies’ denominations tend to grow fast on average but this happens through sudden jumps, as revealed by the fact that their median growth is in fact non positive, the skewness is positive and large and excess kurtosis is massive, similarly to the returns on digital currencies.

Figure 1. Dynamic probability of inclusion of a given predictor for Bitcoin returns. The plots report the recursive estimates of the posterior inclusion probabilities of Bitcoin Google search frequency, the returns on gold and on the SMB portfolio.

3. The dynamics of returns predictability

Before investigating the out-of-sample predictability of cryptocurrencies vis-á-vis traditional asset classes, we first look at the in-sample dynamic exposure of returns to proxy of sources of risk. More specifically, we estimate Equations (3) and (4) for each of the cryptocurrency returns series all the available data. All estimates are based on rather 'uninformative' prior specifications for the hyper-parameters of interest. In a very pragmatic way, we set the initial probability of inclusion at time $t = 0\pi_{0\,|\,0,q} = 1/Q$ and $\hat{\boldsymbol{\theta}}_0(q) = \mathbf{0}$ for all values of $q = 1, 2, \cdots, Q$, with Q the total number of possible model combinations. In addition, we set $\Sigma_{0\,|\,0}(q) = 1000\mathbf{I}_{P+1}$,. This prior is of course, very diffuse and reflects considerable, initial uncertainty on the financial nature of crytpocurrencies as an asset class. Finally, in view of our use of weekly data we set the RiskMetrics-style parameter to $\delta = 0.96$ for the dynamics of the residuals variance $\hat{h}_{t+1}(q)$ to guarantee a not too slow decay. More details on the estimation procedure and the choice of the prior hyper-parameters are provided in Appendix A.

3.1. *Cryptocurrency returns*

The full set of predictors consists of $P = 17$ variables which implies that we have $Q^j = 131{,}072$ possible model combination possibilities.[7] To ease the interpretation of the estimated coefficients, the predictors in $\mathbf{z}_t$ are standardized by dividing their values by their recursively estimated standard deviation, as in Koop and Korobilis (2012). Therefore, we can interpret the coefficients as the cryptocurrency return reaction to a one-standard deviation shock to each predictor, when all the others are held constant. For each asset we report three sets of results: the dynamic probability of inclusion for the three most relevant predictors – those predictors for which the inclusion probability is more than 50% for at least a third of the sample –, the corresponding estimates of the time-varying regression coefficients $\hat{\boldsymbol{\theta}}^{DMA}_{t+1\,|\,t}$, and the unconditional summary statistics for each of the remaining dynamic betas.

3.1.1. *Bitcoin*

With reference to Bitcoin data, in Figure 1, we report the probability of inclusion for Bitcoin google searches, the returns on the first nearby Gold futures contract, and the returns on the size equity risk factor over the period between January 2011 and January 2021. More precisely, the plot shows the time t probability of the predictors to be included across the $Q^{Bitcoin} = 131,072$ models, i.e. $\sum_{l=1}^{Q} \pi_{t+1\,|\,t,l}\, I\{\hat{\theta}_{p,t}(l) \neq \varnothing\}$, where $I\{\hat{\boldsymbol{\theta}}_t(l) \neq \varnothing\}$ indicates that predictor $p = 1, 2, \cdots, P$ is included in the vector $\mathbf{z}_t(l)$.

Two interesting facts emerge; first, in the case of Google searches and of Gold futures returns, we provide evidence of a persistently strong probability of inclusion of these two predictors. Bitcoin returns are predicted by these two variables for a significant fraction of the sample, during 2012–2014 and then again after 2017, when $\Pr(z^j_{p,t}$ included) comes close to 0.6 (60%). Interestingly, the predictive content of the standardized SMB factor is quite erratic over the sample, but nevertheless quite significant especially for the period between early 2016 to early 2020. The unreported dynamics of the other predictors is also quite volatile. Most of the remaining

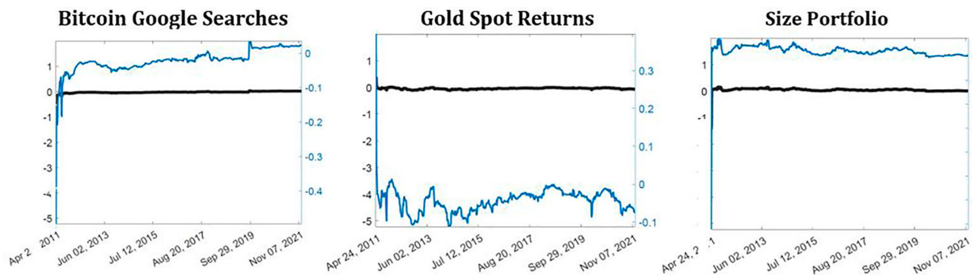

Figure 2. Recursive dynamic model average coefficient estimates for Bitcoin returns. The thicker (black) line reports the recursive estimates on a fixed scale marked on the left axis, the lighter (blue) curve concerns estimates on a variable (right) scale to enhance readability.

predictors are included in the model with posterior probabilities that oscillate between 0.2 and 0.4 over time with few peaks and clear drifts. The probability of inclusion of AMD returns, RMW portfolio, and TSMC stock returns peak between 2011 and 2012 to then decline between 2013 and 2016; the probability for AMD returns, the Chinese electricity index growth rate, the returns on the HML portfolio, the returns on NVIDIA, and silver prices trend up in the last part of the sample, 2018–2020 including the first part of the Covid-19 pandemic crisis. However, we fail to detect particular swings in the most relevant predictors in the last 15 weeks of the sample, characterized by the pandemic shock.

Figure 2 reports the time-varying estimates of the corresponding vector $\hat{\boldsymbol{\theta}}^{DMA}_{t+1\,|\,t}$ for each of the three most relevant predictors. For each estimated dynamic beta coefficient, we report the same values on a two different scales to allow for more meaningful interpretation and commentary. On left-hand axis, we adopt a scale homogeneous across different plots. On the right-hand side, we report the estimated coefficients on a scale that varies across different predictors, for better readability of the variation of the coefficients over time. The evidence suggest that – with minor exceptions concerning for instance Bitcoin momentum between 2017 and 2018, and most importantly the rate of growth of Google searches concerning the word 'Bitcoin' (consistently with the role of attention index in Drobetz, Momtaz, and Schröder 2019; Li and Yi 2019) – there is very limited overall predictability, across model averages, of Bitcoin returns: most of the time and for most predictors (apart from the Bitcoin momentum), the estimated coefficients carry values close to zero. The thinner curves (see the right scale) show that even if the slopes are modest, the coefficients display considerable variation: also because of the small number of observations the DMAs are based on, they oscillate considerably early on in the sample. A visual inspection of the remaining, unreported plots of the recursively estimated vector $\hat{\boldsymbol{\theta}}^{DMA}_{t+1\,|\,t}$ shows that most of the remaining plots flat-line at zero, indicating that the corresponding predictor is never selected in the optimal mix.[8] Overall, there is evidence that especially over the initial and final parts of the sample (2011–2013, and recently 2020), Bitcoin returns are predicted by very simple models, occasionally just featuring a constant, the returns on Gold first-nearby futures, and – over some sub-sample – the rate of growth in the Google searches, which confirms how Bitcoin is largely an asset linked to mass mood swings and attention. For comparison, in Figure B.1 of the Online Appendix, we show that the recursive estimation results of linear predictive regressions obtained from DMA and DMS substantially differ from those one would have naively derived from recursive OLS.[9]

The top panel of Table 2 presents statistics summarizing the information in Figure 2 and the unconditional statistics for the unreported dynamic betas. The means and medians of the recursively estimated DMA slopes are generally modest and very far off the ± 1 one-to-one reactivity. The only partial exception is $\hat{\theta}^{DMA}_{p,t+1\,|\,t}$ for $p =$ SMB, characterized by a mean of 0.043 and by a median of 0.045. Interestingly, as the 'buzz' concerning Bitcoin spreads, the crypto asset may have turned in an alternative asset able to partially replace gold in the portfolios of some investors, which explains the negative albeit small (the mean is -0.044 and the median -0.038) coefficients associated to gold returns, although the marginal impact is rather modest. Interestingly, the growth rate of the Google searches concerning Bitcoin (an obvious clear proxy for investors' attention) turns out to imply on average modest coefficient estimates. More generally–as shown by the bold-facing in the last two columns of the table–for only 2 predictors out of 17, we have evidence of their empirical (probability-weighted across models) 90% confidence band failing to include a zero coefficient: the SMB factor and gold spot returns. In fact, we have also computed the unconditional average number of predictors included by the DMA approach as follows,

$$Size_j^{DMA} \equiv \sum_{t=1}^{T^j-1} \sum_{l=1}^{Q^j} \pi_{t+1\,|\,t,l} Size(\hat{\boldsymbol{\theta}}_t^j(l)) \quad j = \text{BTC, LTC, XRP, ETH} \tag{5}$$

where $Size(\hat{\boldsymbol{\theta}}_t^j(l))$ returns the number of predictors included in model l at time t. It turns out that under DMA, $Size_{Bitcoin}^{DMA} = 5.62$, which appears to be rather modest. Inspection of unreported plots of $Size(\hat{\boldsymbol{\theta}}_t^{Bitcoin}(l))$ shows that the average size of the models favored by the recursively updated probabilities somewhat increases (to exceed 6) between 2017 and 2019 but declines again in the final part of the sample. However, there are long spells (2011 and 2013) in which $Size(\hat{\boldsymbol{\theta}}_t^{Bitcoin}(l))$ falls below three: almost nothing but Bitcoin momentum and some precious metal returns forecast Bitcoin returns. The bottom panel of Table 2 summarizes the key DMS results: the means and medians of the standardized predictors are generally small, never exceed 0.1 in absolute value, while their (across time series and cross-sections of models) empirical 90 % confidence interval always includes zero. In other words, even though some of the predictors come close to display some predictive power for Bitcoin returns, this is generally muted in a DMS framework.[10]

3.1.2. Litecoin

Figure 3 reports the probability of inclusion for LTC google searches, the lagged returns on BTC, and the excess returns on the aggregate equity market for developed economies. With the exception of the Litecoin Google search growth rate, the probabilities of models including any predictors are generally modest, seldom exceeding 0.50; moreover, most plots show declining probabilities over time, after starting out at a relatively high level between 2013 and 2015. The plot for the probability of inclusion of the rate of growth of the Google searches describes an inverse ‘V-shaped’ evolution, starting out at 1, declining by late 2016 to almost zero, and then bouncing back to 1 by the end of 2017. The other unreported probabilities of inclusion are generally modest and oscillate between 0.20 and 0.40. These results points towards a weaker predictability of LTC returns on a weekly basis.

Figure 4 reports the posterior estimates of the dynamic betas for the same three predictors. The statistics are based on weekly data, expressed in US dollars. The percentile statistics are across time and probability-weighted within time periods. We have boldfaced pairs of 5th–95th empirical percentile statistics that fail to include zero, which implies that in overall terms (under the assumption of independence over time), the series of recursively estimated slopes are statistically significant at a 10% size. The slope coefficient projecting Litecoin returns on the rate of growth of Litecoin Google searches shows a generally declining coefficient, also characterized by a partial ‘V-shaped’ evolution, in which the coefficient declines from 1 to almost zero between 2013 and late 2016, to then climb back up towards 0.4 from late 2017. On the contrary, the CAPM-style DMA betas for excess market returns appear to steadily decline over time, from coefficients of approximately 0.3 in 2013 to zero or even slightly negative values by the end of 2017. A local, predictive CAPM in which Bitcoin returns explain Litecoin returns does not seem to perform very well in spite the predictor is included in the model average quite often, as shown in Figure 3. The logic of inclusion of lagged Bitcoin returns in the prediction of other crypto assets is that Bitcoin is susceptible to sentiment trading because it is the most traded cryptocurrency. Thus, it can capture periods when trading forces that are unrelated to fundamentals are strong. Also, Bitcoin is the largest cryptocurrency in terms of market capitalization and it can be used as a proxy for digital currency market-wide risk. Finally, especially on minor trading venues for other cryptocurrencies, Bitcoin is routinely used to pay transactions in other assets. Whether Bitcoin captures investor sentiment or is a proxy for systematic cryptocurrency risk, it is important that we include it in our predictive models. In any event, because Bitcoin itself is mostly explained by the rate of growth in the Google searches of the word Bitcoin and now it enters as a predictor of Litecoin returns, it seems that also Litecoin is eventually an almost entirely attention-driven asset. Similar to BTC returns, Table 3 shows that the associated coefficients are generally small and hardly economically significant. The table shows at a glance–especially when we focus on whether the empirical 90% confidence bands for the estimated

Table 2. Summary statistics for recursively estimated DMA and DMS predictive regression coefficients: Bitcoin returns.

	Mean	Median	St. Dev.	Min	Max	5th percentile	95th percentile
				Dynamic Model Average Coefficients			
Bitcoin Lagged returns	0.0131	0.0132	0.0279	0.1284	−0.0627	−0.0328	0.0686
Bitcoin Momentum	−0.0109	−0.0020	0.0346	0.1030	−0.1943	−0.0790	0.0228
Developed Markets portfolio returns	0.0247	0.0190	0.0301	0.3990	−0.0689	−0.0076	0.0691
SMB portfolio returns	0.0456	0.0445	0.0708	0.1516	−1.4460	**0.0040**	**0.1007**
HML portfolio returns	−0.0304	0.0177	0.2660	0.1252	−5.2172	−0.2960	0.0920
RMW portfolio returns	−0.0101	0.0000	0.0765	0.0785	−0.8339	−0.0790	0.0531
CMA portfolio returns	−0.0442	−0.0138	0.1125	1.5847	−0.7310	−0.1582	0.0153
Momentum portfolio returns	−0.0208	−0.0002	0.0654	0.0902	−0.6324	−0.1144	0.0389
Gold spot returns	−0.0436	−0.0380	0.0325	0.3957	−0.1131	**−0.0940**	**−0.0064**
Platinum spot returns	0.0366	0.0314	0.0488	0.4940	−0.0167	−0.0085	0.0969
Silver spot returns	−0.0168	−0.0078	0.0713	0.0773	−1.1923	−0.1171	0.0479
Growth of Google Bitcoin searches	−0.0217	−0.0211	0.0361	0.0343	−0.4961	−0.0648	0.0205
VW returns on Chinese electricity stocks	0.0326	−0.0014	0.1186	1.9902	−0.0551	−0.0402	0.1322
VW returns on US electricity stocks	−0.0346	−0.0099	0.1136	0.0585	−1.0062	−0.1494	0.0208
NVIDIA individual stock returns	−0.0024	0.0001	0.0344	0.0497	−0.1496	−0.0795	0.0413
AMD individual stock returns	−0.0230	0.0001	0.0639	0.0294	−0.3485	−0.1957	0.0229
TSMC individual stock returns	0.0237	0.0081	0.0646	0.7691	−0.0465	−0.0304	0.1739
Residual predictive variance (ht)	1.1225	0.7088	0.9799	4.6949	0.2069	**0.2777**	**3.5231**
				Dynamic Model Selection Coefficients			
Bitcoin Lagged returns	0.0179	0.0000	0.0404	0.1858	−0.0145	−0.0001	0.1269
Bitcoin Momentum	−0.0092	0.0000	0.0605	0.0000	−0.7103	**0.0000**	**0.0000**
Developed Markets portfolio returns	0.0332	0.0000	0.0951	0.5202	0.0000	**0.0000**	**0.2867**
SMB portfolio returns	0.0848	0.0694	0.0973	0.4712	0.0000	**0.0000**	**0.2734**
HML portfolio returns	−0.0222	0.0000	0.2966	0.1459	−3.2369	0.0000	0.0000
RMW portfolio returns	−0.0078	0.0000	0.1431	0.2884	−1.3682	0.0000	0.0000
CMA portfolio returns	−0.0352	0.0000	0.1593	0.0000	−1.9472	**−0.2877**	**0.0000**
Momentum portfolio returns	−0.0125	0.0000	0.1316	0.0000	−1.4274	0.0000	0.0000
Gold spot returns	−0.0123	0.0000	0.0406	0.0000	−0.2090	**−0.1093**	**0.0000**
Platinum spot returns	0.0171	0.0000	0.0614	0.6741	0.0000	**0.0000**	**0.1467**
Silver spot returns	−0.0081	0.0000	0.0289	0.1622	−0.1204	**−0.0912**	**0.0000**
Growth of Google Bitcoin searches	−0.0244	0.0000	0.0503	0.0000	−0.3544	**−0.1085**	**0.0000**
VW returns on Chinese electricity stocks	0.0292	0.0000	0.1302	1.5914	−0.1087	**0.0000**	**0.2316**
VW returns on US electricity stocks	−0.0213	0.0000	0.1631	0.0000	−1.6153	0.0000	0.0000
NVIDIA individual stock returns	0.0080	0.0000	0.0512	0.1043	−0.4499	**0.0000**	**0.0776**
AMD individual stock returns	−0.0248	0.0000	0.1270	0.0000	−0.7794	0.0000	0.0000
TSMC individual stock returns	0.0085	0.0000	0.0737	0.7216	−0.0727	0.0000	0.0000
Residual predictive variance (ht)	1.0427	0.6785	0.9199	4.4123	0.1935	**0.2603**	**3.2107**

Notes: The statistics are based on weekly data, expressed in US dollars. The percentile statistics are across time and probability-weighted within time periods. We have boldfaced pairs of 5th–95th empirical percentile statistics that fail to include zero, which implies that in overall terms (under the assumption of independence over time), the series of recursively estimated slopes are statistically significant at a 10% size.

coefficients happen to include zero or not–that there is even less predictability in the case of Litecoin vs. Bitcoin. With the only exception of the dynamic beta on the Google SVI, i.e. $\hat{\theta}^{DMA}_{t+1\,|\,t(SVI)}$, none of the other betas seem to be unconditionally significant at a standard 90% confidence level.

3.1.3. Ripple

Table 4 reports summary statistics for the dynamic betas $\hat{\boldsymbol{\theta}}^{DMA}_{t+1\,|\,t}$ for the April 2015–December 2021 sample. The bottom panel of the table shows that virtually none of the predictors significantly enters the model with posterior median probabilities which are all eventually zero. In fact, like in the case of Litecoin, Ripple returns are not forecast by lagged Bitcoin returns, in fact with a sample median probability of inclusion close to zero. Also, in this case, investors' attention in the form of Google searches is the key variable to predict Ripple returns. To summarize, Table 4 shows that none of the predictors consistently enters the best selected model so to end up delivering mean and median coefficient estimates that are essentially zero. In fact, we find that $Size^{DMS}_{Ripple} = 0.5$, which means that on average the best selected model just includes the constant, i.e. no predictability exists. The absence of significant predictability is confirmed by the time series of posterior inclusion probabilities and the

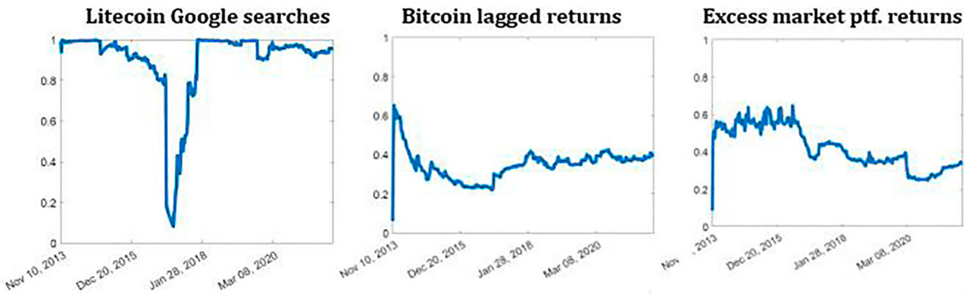

Figure 3. Dynamic probability of inclusion of a given predictor for Litecoin returns. The plots report the recursive estimates of the posterior inclusion probabilities of Litecoin Google search frequency, Bitcoin lagged returns, and the aggregate value-weighted market index.

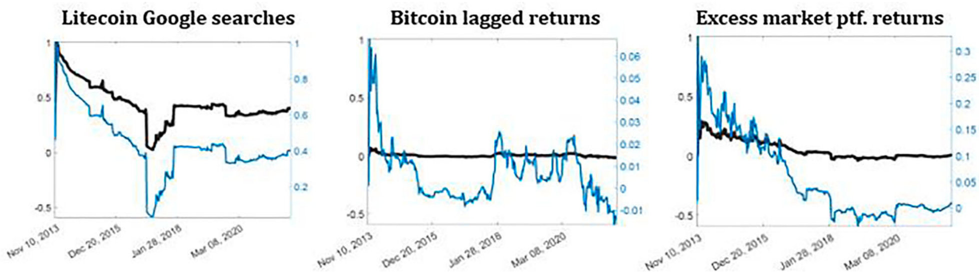

Figure 4. Recursive dynamic model average coefficient estimates for Litecoin returns. The thicker (black) line reports the recursive estimates on a fixed scale marked on the left axis, the lighter (blue) curve concerns estimates on a variable (right) scale to enhance readability.

dynamic beta estimates. For the sake of brevity, and in order to streamline the discussion, we decided to not reporting those estimates given they are virtually always not statistically significant.[11]

3.1.4. Ethereum

The empirical evidence concerning Ethereum stands out as a bit dissimilar compared to our findings for the other cryptocurrency returns. First, as showed in Figure 5, concerning the posterior probability of inclusion of the investment factor lagged returns, ETH Google search, and the returns on the silver first nearby futures contract. In the case of CMA returns, the probability starts out at zero but rapidly increases reaching a maximum in mid-2017, at almost 0.9; the probability subsequently declines rather smoothly to 0.2. The plot of silver returns is representative of the remaining 14 plots that we have not reported to save space: the probability of inclusion is generally not very high but it also gyrates wildly, occasionally spiking up to exceed 0.4. Similarly, the probability of inclusion of the Google attention predictor shows a fast increase in late 2016, after which it stabilizes at a level of 0.5−0.6, which is however less than LTC.

Figure 6 reports the time-varying betas estimates. Interestingly, the plot for lagged investment factor returns describes a shape that mirrors in reverse the probability of inclusion of CMA returns, which makes the contribution to predicted Ethereum returns rather steady. On the opposite, the DMA coefficient for the ETH SVI spikes in a manner consistent with its probability of inclusion, which implies that for a few weeks in late 2017 the full predictive power came from this predictor. Finally, the DMA coefficient for silver spot returns is rather volatile, even though the tight (right) scale should be taken into account.

Table 5 shows that, although means and medians of the recursively estimated DMA coefficients are generally small, the empirical 90 % bands do not show widespread evidence of resilient predictability. As a matter of fact, with the only partial exception of the returns on the silver futures returns, for which the mean (median) coefficient is 0.063 (0.044), none of the other predictors are consistently positive or negative throughout the sample, as shown by the 90% credibility intervals. Therefore, there is some predictability for this crypto asset with $Size^{DMA}_{Ethereum} = 4.65$, which is fairly high. Finally, the DMS panel of Table 5 shows that only four models

Table 3. Summary statistics for recursively estimated DMA and DMS predictive regression coefficients: Litecoin returns.

	Mean	Median	St. Dev.	Min	Max	5th percentile	95th percentile
				Dynamic Model Average Coefficients			
Bitcoin Lagged returns	0.0066	0.0041	0.0127	0.0678	−0.0167	−0.0075	0.0254
Litecoin Momentum	−0.0053	−0.0038	0.0389	0.2208	−0.5928	−0.0168	0.0123
Developed Markets portfolio returns	0.0572	0.0248	0.0792	0.3282	−0.0355	−0.0204	0.2103
SMB portfolio returns	−0.0048	−0.0090	0.0278	0.3063	−0.0620	−0.0365	0.0211
HML portfolio returns	0.0100	0.0067	0.0304	0.3623	−0.0519	−0.0305	0.0634
RMW portfolio returns	−0.0155	0.0014	0.0604	0.0241	−0.5692	−0.0913	0.0176
CMA portfolio returns	0.0086	0.0028	0.0400	0.1586	−0.1636	−0.0366	0.0937
Momentum portfolio returns	−0.0219	−0.0136	0.0323	0.1529	−0.1421	−0.0993	0.0054
Gold spot returns	−0.0010	−0.0020	0.0234	0.0616	−0.1791	−0.0349	0.0387
Platinum spot returns	−0.0202	−0.0206	0.0749	0.1630	−0.2881	−0.1322	0.1401
Silver spot returns	0.0032	0.0075	0.0320	0.1033	−0.1063	−0.0835	0.0479
Growth of Google Litecoin searches	0.4396	0.4175	0.1763	1.0073	0.0263	**0.1412**	**0.7828**
VW returns on Chinese electricity stocks	−0.0139	−0.0113	0.0369	0.0423	−0.3678	−0.0682	0.0264
VW returns on US electricity stocks	−0.0066	−0.0046	0.0191	0.0375	−0.1962	−0.0404	0.0155
NVIDIA individual stock returns	−0.0290	−0.0070	0.0622	0.0651	−0.3357	−0.1421	0.0263
AMD individual stock returns	−0.0376	−0.0061	0.0854	0.0306	−0.5023	−0.2346	0.0216
TSMC individual stock returns	0.0021	0.0017	0.0106	0.0835	−0.0168	−0.0105	0.0148
Residual predictive variance (ht)	0.7621	0.5799	0.5625	3.6710	0.1798	**0.2738**	**1.9329**
				Dynamic Model Selection Coefficients			
Bitcoin Lagged returns	0.0006	0.0000	0.0114	0.2347	0.0000	0.0000	0.0000
Litecoin Momentum	−0.0019	0.0000	0.0390	0.0000	−0.8047	0.0000	0.0000
Developed Markets portfolio returns	0.0123	0.0000	0.0692	1.0199	0.0000	0.0000	0.0000
SMB portfolio returns	0.0024	0.0000	0.0496	1.0224	0.0000	0.0000	0.0000
HML portfolio returns	—	—	—	—	—	—	—
RMW portfolio returns	−0.0033	0.0000	0.0509	0.0000	−0.9394	0.0000	0.0000
CMA portfolio returns	0.0028	0.0000	0.0283	0.3100	0.0000	0.0000	0.0000
Momentum portfolio returns	—	—	—	—	—	—	—
Gold spot returns	−0.0022	0.0000	0.0444	0.0000	−0.9159	0.0000	0.0000
Platinum spot returns	0.0195	0.0000	0.0653	0.1985	−0.4718	0.0000	0.1795
Silver spot returns	0.0131	0.0000	0.0454	0.5848	0.0000	0.0000	0.1197
Growth of Google Litecoin searches	0.4496	0.4190	0.1575	1.0744	0.0000	**0.3029**	**0.6955**
VW returns on Chinese electricity stocks	−0.0031	0.0000	0.0635	0.0000	−1.3082	0.0000	0.0000
VW returns on US electricity stocks	−0.0012	0.0000	0.0249	0.0000	−0.5125	0.0000	0.0000
NVIDIA individual stock returns	—	—	—	—	—	—	—
AMD individual stock returns	−0.0275	0.0000	0.1284	0.0000	-1.2459	0.0000	0.0000
TSMC individual stock returns	0.0009	0.0000	0.0185	0.3804	0.0000	0.0000	0.0000
Residual predictive variance (ht)	0.6969	0.5500	0.4485	2.5590	0.0750	**0.2564**	**1.5948**

Notes: The statistics are based on weekly data, expressed in US dollars. The percentile statistics are across time and probability-weighted within time periods. We have boldfaced pairs of 5th–95th empirical percentile statistics that fail to include zero, which implies that in overall terms (under the assumption of independence over time), the series of recursively estimated slopes are statistically significant at a 10% size.

are ever selected as the best one for forecasting purposes, while the remaining never get picked. Of course, this also reflects the ‘spiky’ dynamics of some of the plots in Figures 5 and 6. Consistently with this remark, we find $Size^{DMS}_{Ethereum} = 1.26$.

3.2. Are cryptocurrencies just another currency?

An important, additional test is based on assessing whether cryptocurrencies are eventually just an alternative asset class that differ from standard fiat currencies. Therefore, the specialty of cryptocurrencies and possibly their segmentation from remaining currencies can be best established by a comparison with the predictability patterns revealed by the US dollar exchange rate. We obtain data on the US dollar effective exchange rate (i.e. the trade-weighted exchange rate of the US dollar against a subset of major currencies including the Euro, Canada, Japan, United Kingdom, Switzerland, Australia, and Sweden) for a January 2011 to December 2021 sample, similarly to Bitcoin data. The logic is that in the same way in which digital currency returns measure relative percentage changes in the value of a cryptocurrency vis-a-vis the US dollar, the log changes in the effective nominal exchange rate of the US dollar measure the relative changes of its value vis-a-vis foreign currencies.

Table 4. Summary statistics for recursively estimated DMA and DMS predictive regression coefficients: Ripple returns.

	Mean	Median	St. Dev.	Min	Max	5th percentile	95th percentile
	Dynamic Model Average Coefficients						
Bitcoin Lagged returns	−0.0157	−0.0113	0.0133	0.0032	−0.0484	**−0.0388**	**−0.0007**
Ripple Momentum	−0.0180	−0.0058	0.0310	0.0245	−0.1311	−0.0860	0.0087
VW CRSP Mkt excess returns	0.0038	0.0037	0.0053	0.0305	−0.0115	−0.0023	0.0119
SMB portfolio returns	0.0124	0.0063	0.0136	0.0555	−0.0047	**−0.0027**	**0.0359**
HML portfolio returns	−0.0021	−0.0025	0.0104	0.0790	−0.0158	−0.0148	0.0128
RMW portfolio returns	0.0131	0.0125	0.0112	0.0971	−0.0029	**0.0006**	**0.0267**
CMA portfolio returns	−0.0068	−0.0059	0.0124	0.0156	−0.1078	−0.0109	0.0008
Momentum portfolio returns	0.0045	0.0018	0.0148	0.0982	−0.0262	−0.0203	0.0221
Gold spot returns	0.0073	0.0070	0.0086	0.0471	−0.0070	−0.0047	0.0236
Platinum spot returns	−0.0070	−0.0047	0.0168	0.0241	−0.0631	−0.0337	0.0186
Silver spot returns	−0.0018	−0.0002	0.0157	0.0166	−0.1313	−0.0092	0.0084
Growth of Google Ripple searches	0.2098	0.2784	0.1611	0.3906	−0.1793	−0.1002	0.3711
VW returns on Chinese electricity stocks	−0.0047	−0.0041	0.0072	0.0160	−0.0157	−0.0143	0.0092
VW returns on US electricity stocks	0.0028	0.0035	0.0052	0.0209	−0.0100	−0.0038	0.0114
NVIDIA individual stock returns	−0.0043	−0.0029	0.0068	0.0081	−0.0410	−0.0174	0.0042
AMD individual stock returns	0.0010	0.0002	0.0060	0.0416	−0.0074	−0.0066	0.0123
TSMC individual stock returns	0.0121	0.0076	0.0132	0.0478	−0.0074	−0.0031	0.0360
Residual predictive variance (ht)	0.9659	0.3632	1.4479	6.5742	0.0572	**0.0658**	**4.6106**
	Dynamic Model Selection Coefficients						
Bitcoin Lagged returns	−0.0006	0.0000	0.0068	0.0000	−0.0747	0.0000	0.0000
Ripple Momentum	−0.0264	0.0000	0.0630	0.0000	−0.2171	−0.1784	0.0000
VW CRSP Mkt excess returns	—	—	—	—	—	—	—
SMB portfolio returns	—	—	—	—	—	—	—
HML portfolio returns	—	—	—	—	—	—	—
RMW portfolio returns	—	—	—	—	—	—	—
CMA portfolio returns	—	—	—	—	—	—	—
Momentum portfolio returns	−0.0002	0.0000	0.0030	0.0000	−0.0409	0.0000	0.0000
Gold spot returns	0.0024	0.0000	0.0110	0.0575	0.0000	0.0000	0.0000
Platinum spot returns	—	—	—	—	—	—	—
Silver spot returns	—	—	—	—	—	—	—
Growth of Google Ripple searches	0.2042	0.2839	0.1591	0.4204	0.0000	**0.0000**	**0.3618**
VW returns on Chinese electricity stocks	—	—	—	—	—	—	—
VW returns on US electricity stocks	—	—	—	—	—	—	—
NVIDIA individual stock returns	—	—	—	—	—	—	—
AMD individual stock returns	—	—	—	—	—	—	—
TSMC individual stock returns	—	—	—	—	—	—	—
Residual predictive variance (ht)	0.9441	0.3329	1.4474	6.5541	0.0540	**0.0588**	**4.6158**

Notes: The statistics are based on weekly data, expressed in US dollars. The percentile statistics are across time and probability-weighted within time periods. We have boldfaced pairs of 5th–95th empirical percentile statistics that fail to include zero, which implies that in overall terms (under the assumption of independence over time), the series of recursively estimated slopes are statistically significant at a 10% size.

In fact, the comparison of cryptocurrencies with traditional fiat currencies and exchange rates has a few interesting antecedents in the literature, see, e.g. Baur, Hong, and Lee (2018) and Urom et al. (2020). Although in Table 1 we may notice that the 0.06 Sharpe ratio characterizing a long-short position in US dollars vs. a basket of other currencies is inferior to that typical of cryptocurrencies, Table 6 shows that – even using predictor variables that are hardly optimized to be consistent with the exchange rate literature – there is considerably more predictability in US dollar log-changes vs. what we have reported for cryptocurrencies. In Table 6, six predictors out of seventeen are characterized by empirical 90% confidence interval for the DMA coefficients that fail to contain zero; in the case of DMS estimates, 15 coefficients give signs of accurate and persistent inclusion in the predictive relationship. In unreported plots (available upon request) there is visible evidence, especially in comparison to Figures 1–6, of some forecasting power of our selected predictors for the trade-weighted value of the US dollar. In particular, the US Market, SMB, and HML portfolio excess returns, platinum and silver spot returns, and the returns on US electricity stocks are all characterized by time-varying coefficients that seem to capture pockets of predictability.[12]

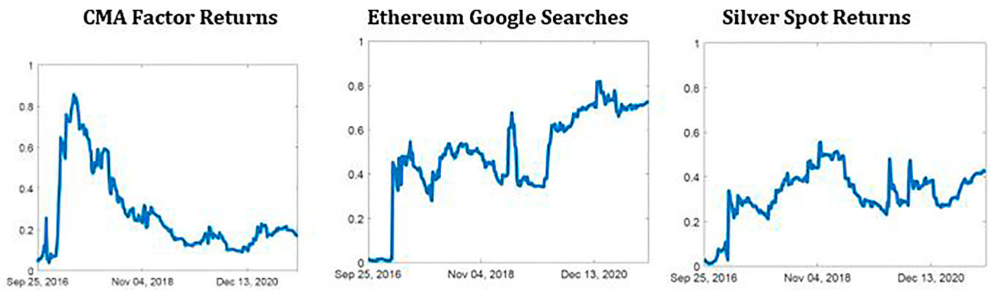

Figure 5. Dynamic probability of inclusion of a given predictor for Ethereum returns. The thicker line reports the recursive estimates of the posterior inclusion probabilities of the long-short Investment factor portfolio, Ethereum Google searches, and silver spot returns.

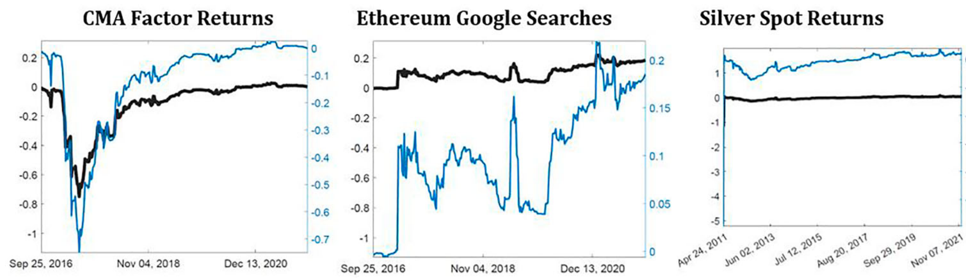

Figure 6. Recursive dynamic model average coefficient estimates for Ethereum returns. The thicker (black) line reports estimates on a fixed scale marked on the left axis, the lighter (blue) curve concerns estimates on a variable (right) scale to enhance readability.

4. Forecasting comparison and portfolio performance

One of the key questions of the paper is whether the empirical evidence of predictability (or lack thereof) in the returns of cryptocurrencies may justify a claim that these are segmented from more traditional asset classes. To tackle this question we test whether DMA and DMS models imply an overall amount of out-of-sample (OOS) predictability for cryptocurrencies that is significantly different from that recorded for traditional asset classes. We test this assumption both through standard loss functions, such as the squared loss that leads us to rank alternative models and to quantify OOS predictability using MSFE criteria, as well as more economically grounded loss functions. More specifically, we perform a simple and yet robust recursive expected utility-based asset allocation exercise to test whether the economic value generated from the predictability in digital currency returns may compare to that recorded for other, traditional asset classes. Importantly, we depart from earlier literature (see, e.g. Chuen et al. 2017; Cheah et al. 2022) and do not simply assume a mean-variance objective but instead work under constant relative risk aversion (CRRA), power expected utility that penalizes also negative asymmetries and fat tails in the (predicted) densities of asset returns.

4.1. Realised OOS predictive performance

We systematically compute and record the OOS realized forecasting performance of alternative models for different asset classes, with the differential performance across digital currencies and other asset classes in the spot light. Table 7 provides the crucial piece of evidence in this paper and it is organized as follows. The top panel concerns crypto returns and, for comparison, the bottom panel the other asset classes investigated so far, including international and US stocks and gold. In this regard, we caution against the heterogeneity of the samples available, even though all recursive OOS forecasts for the benchmarks have been obtained on a sample identical to that for Bitcoin, July 2011 to December 2021. This implies that strictly speaking, the results are perfectly comparable only when Bitcoin is involved. All models are initialized on a 6-month weekly sample of observations. Table 7 has four blocks of results concerning four alternative prediction strategies: naive recursive OLS; recursive

Table 5. Summary statistics for recursively estimated DMA and DMS predictive regression coefficients: Ethereum returns.

	Mean	Median	St. Dev.	Min	Max	5th percentile	95th percentile
				Dynamic Model Average Coefficients			
Bitcoin Lagged returns	−0.0076	0.0030	0.0512	0.0578	−0.2539	−0.1212	0.0355
Ethereum Momentum	−0.0268	0.0256	0.1907	0.1045	−1.1300	−0.4777	0.0552
VW CRSP Mkt excess returns	0.0264	0.0094	0.0400	0.1268	−0.0189	−0.0110	0.1067
SMB portfolio returns	0.0137	0.0128	0.0221	0.1503	−0.0267	−0.0095	0.0301
HML portfolio returns	−0.0307	−0.0038	0.0579	0.0265	−0.2682	−0.1644	0.0153
RMW portfolio returns	−0.0060	0.0010	0.0303	0.0436	−0.1487	−0.0627	0.0231
CMA portfolio returns	−0.1184	−0.0340	0.1726	0.0254	−0.7485	−0.5359	0.0116
Momentum portfolio returns	−0.0700	−0.0200	0.0898	0.0342	−0.3145	−0.2423	0.0110
Gold spot returns	0.0456	0.0371	0.0362	0.1287	−0.0171	0.0038	0.1087
Platinum spot returns	0.0522	0.0373	0.0558	0.2139	−0.0534	−0.0171	0.1525
Silver spot returns	0.0634	0.0444	0.0483	0.2122	0.0037	**0.0108**	**0.1657**
Growth of Google Ethereum searches	0.0985	0.0961	0.0548	0.2193	−0.0065	−0.0024	0.1802
VW returns on Chinese electricity stocks	−0.0336	−0.0079	0.0578	0.0164	−0.2681	−0.1659	0.0060
VW returns on US electricity stocks	0.0012	0.0026	0.0319	0.1146	−0.0646	−0.0459	0.0554
NVIDIA individual stock returns	0.0293	0.0269	0.0259	0.1168	−0.0697	−0.0007	0.0749
AMD individual stock returns	−0.0157	−0.0162	0.0135	0.0085	−0.0507	−0.0389	0.0053
TSMC individual stock returns	0.0136	0.0050	0.0315	0.1365	−0.0699	−0.0298	0.0711
Residual predictive variance (ht)	1.1500	0.8638	0.6575	2.8409	0.3111	**0.4116**	**2.4925**
				Dynamic Model Selection Coefficients			
Bitcoin Lagged returns	—	—	—	—	—	—	—
Ethereum Momentum	−0.0329	0.0000	0.2232	0.2789	−1.2785	0.0000	0.0000
VW CRSP Mkt excess returns	—	—	—	—	—	—	—
SMB portfolio returns	0.0067	0.0000	0.0492	0.3728	0.0000	0.0000	0.0000
HML portfolio returns	−0.0037	0.0000	0.0431	0.0000	−0.5157	0.0000	0.0000
RMW portfolio returns	—	—	—	—	—	—	—
CMA portfolio returns	−0.0861	0.0000	0.1999	0.0000	−0.5979	**−0.5659**	**0.0000**
Momentum portfolio returns	—	—	—	—	—	—	—
Gold spot returns	—	—	—	—	—	—	—
Platinum spot returns	0.0467	0.0000	0.0905	0.2594	0.0000	**0.0000**	**0.2329**
Silver spot returns	0.0274	0.0000	0.0826	0.3053	0.0000	**0.0000**	**0.2883**
Growth of Google Ethereum searches	0.1309	0.2044	0.1157	0.2719	0.0000	**0.0000**	**0.2622**
VW returns on Chinese electricity stocks	—	—	—	—	—	—	—
VW returns on US electricity stocks	—	—	—	—	—	—	—
NVIDIA individual stock returns	—	—	—	—	—	—	—
AMD individual stock returns	—	—	—	—	—	—	—
TSMC individual stock returns	—	—	—	—	—	—	—
Residual predictive variance (ht)	0.9501	0.7862	0.4645	1.9996	0.2089	**0.3144**	**1.7789**

Notes: The statistics are based on weekly data, expressed in US dollars. The percentile statistics are across time and probability-weighted within time periods. We have boldfaced pairs of 5th–95th empirical percentile statistics that fail to include zero, which implies that in overall terms (under the assumption of independence over time), the series of recursively estimated slopes are statistically significant at a 10% size.

OLS based on the full-sample best individual predictor, that is also listed in the table; DMA; DMS. Notice that recursive OLS forecasts based on the full-sample best individual predictor are not true OOS forecasts because they could not have been implemented in real time, as they are based on a selection of predictor that becomes available only as of May 2020. We have purposefully specified this unfeasible predictive model to provide a quantitative indication of what is the best predictive power of simple OLS with our data. Unsurprisingly, in the case of three of the four cryptocurrencies, the best OLS predictor is the rate of growth in the crypto-specific Google searches.

For each model, we report two measures of OOS predictive accuracy, i.e. the root mean-squared forecast error (RMSFE) defined in the conventional way and the OOS R-square, defined as:

$$R^2_{OOS}(j,\mathcal{M}) \equiv 1 - \frac{\sum_{t=P_j+1}^{T}(r_{j,t+1} - \hat{r}^{\mathcal{M}}_{j,t+1\,|\,t})^2}{\sum_{t=P_j+1}^{T}(r_{j,t+1} - \bar{r}_{j,t+1\,|\,t})^2} = 1 - \frac{MSFE(j,\mathcal{M})}{MSFE(j,mean)}, \tag{6}$$

Table 6. Summary statistics for recursively estimated DMA and DMS predictive regression coefficients: returns on the nominal broad U.S. dollar index.

	Mean	Median	St. Dev.	Min	Max	5th percentile	95th percentile
	Dynamic Model Average Coefficients						
Bitcoin Lagged returns	−0.0382	−0.0287	0.0283	−0.0033	−0.1338	**−0.0900**	**−0.0076**
Bitcoin Momentum	−0.0386	−0.0336	0.0391	0.0284	−0.1378	−0.1132	0.0198
VW CRSP Mkt excess returns	0.0249	0.0204	0.0290	0.1458	−0.0503	−0.0206	0.0737
SMB portfolio returns	0.0215	0.0120	0.0333	0.2081	−0.0274	−0.0115	0.0828
HML portfolio returns	−0.1098	−0.0622	0.1074	0.0604	−0.4344	**−0.3642**	**−0.0218**
RMW portfolio returns	0.0284	0.0181	0.0466	0.2055	−0.0489	−0.0317	0.1204
CMA portfolio returns	0.0177	0.0392	0.0805	0.1767	−0.2180	−0.1171	0.1176
Momentum portfolio returns	0.0114	0.0085	0.0247	0.0903	−0.0737	−0.0246	0.0639
Gold spot returns	−0.1506	−0.1136	0.1108	−0.0058	−0.5976	**−0.4067**	**−0.0243**
Platinum spot returns	−0.1592	−0.1920	0.1055	0.0247	−0.3279	**−0.2938**	**−0.0015**
Silver spot returns	−0.2552	−0.2522	0.1074	−0.0205	−0.4789	**−0.4503**	**−0.0621**
Growth of Google Bitcoin searches	0.0138	0.0116	0.0237	0.0865	−0.0531	−0.0191	0.0554
VW returns on Chinese electricity stocks	−0.0296	−0.0314	0.0264	0.0275	−0.1044	−0.0784	0.0137
VW returns on US electricity stocks	−0.1925	−0.1225	0.1708	0.0033	−0.7061	**−0.5491**	**−0.0358**
NVIDIA individual stock returns	−0.0468	−0.0220	0.0584	0.0316	−0.2018	−0.1574	0.0221
AMD individual stock returns	−0.0085	−0.0112	0.0294	0.1012	−0.0810	−0.0493	0.0449
TSMC individual stock returns	0.0265	0.0200	0.0377	0.2553	−0.0241	−0.0122	0.1067
Residual predictive variance (ht)	0.6764	0.6236	0.2020	1.3079	0.3593	**0.4124**	**1.0930**
	Dynamic Model Selection Coefficients						
Bitcoin Lagged returns	−0.0350	0.0000	0.0559	0.0000	−0.2678	**−0.1314**	**0.0000**
Bitcoin Momentum	−0.0374	0.0000	0.0618	0.0618	−0.2170	**−0.1608**	**0.0000**
VW CRSP Mkt excess returns	0.0239	0.0000	0.0557	0.2112	0.0000	**0.0000**	**0.1749**
SMB portfolio returns	0.0144	0.0000	0.0421	0.1989	0.0000	**0.0000**	**0.1221**
HML portfolio returns	−0.1250	0.0000	0.2161	0.0000	−0.9379	**−0.5562**	**0.0000**
RMW portfolio returns	0.0064	0.0000	0.0417	0.3047	0.0000	0.0000	0.0000
CMA portfolio returns	0.0060	0.0000	0.1382	0.3299	−0.4006	−0.2581	0.2517
Momentum portfolio returns	0.0042	0.0000	0.0261	0.2248	0.0000	0.0000	0.0000
Gold spot returns	−0.1135	0.0000	0.1770	0.0000	−0.6802	**−0.4842**	**0.0000**
Platinum spot returns	−0.1699	−0.2126	0.1297	0.0000	−0.4770	**−0.3313**	**0.0000**
Silver spot returns	−0.3089	−0.3219	0.1763	0.0000	−0.5848	**−0.5726**	**0.0000**
Growth of Google Bitcoin searches	0.0109	0.0000	0.0357	0.1591	−0.1229	**0.0000**	**0.1051**
VW returns on Chinese electricity stocks	−0.0302	0.0000	0.0493	0.0000	−0.1758	**−0.1195**	**0.0000**
VW returns on US electricity stocks	−0.2142	−0.1724	0.2012	0.0000	−0.9732	**−0.6398**	**0.0000**
NVIDIA individual stock returns	−0.0586	0.0000	0.0732	0.0283	−0.3045	**−0.1674**	**0.0000**
AMD individual stock returns	−0.0096	0.0000	0.0296	0.0000	−0.1410	**−0.0980**	**0.0000**
TSMC individual stock returns	0.0133	0.0000	0.0512	0.5165	0.0000	**0.0000**	**0.1423**
Residual predictive variance (ht)	0.6096	0.5598	0.1780	1.2184	0.3272	**0.3796**	**0.9683**

Notes: The statistics are based on weekly data. The percentile statistics are across time and probability-weighted within time periods. We have boldfaced pairs of 5th–95th empirical percentile statistics that fail to include zero, which implies that in overall terms (under the assumption of independence over time), the series of recursively estimated slopes are statistically significant at a 10% size.

where j refers to the asset class or cryptocurrency investigated, $\mathcal{M}$ is one of the four predictive frameworks under consideration, and P_j denotes the end of the estimation sample for each potential j. Clearly, $R^2_{OOS}(j, \mathcal{M}) \lesseqqgtr$ according to whether $MSFE(j, \mathcal{M}) \gtreqqless MSFE(j, mean)$ so that a negative $R^2_{OOS}(j, \mathcal{M})$ is not only possible, but also highly meaningful: $R^2_{OOS}(j, \mathcal{M}) < 0$ occurs when a given forecasting model cannot manage to outperform the sample mean based on the naive expanding sample forecast. Of course, compared to the (square root of the) MSFE for a model, $R^2_{OOS}(j, \mathcal{M})$ can be more informative by providing a signed, relative measure. Finally, in Table 7, for each currency or asset class, we have boldfaced the model with the highest $R^2_{OOS}(j, \mathcal{M})$ and the lowest RMSFE.

The message of Table 7 is stark: with no exceptions, all cryptocurrencies – and markedly, Bitcoin among them – are considerably less predictable than other, more traditional asset classes are. While on the one hand this suggests great caution when cryptocurrency investments are approached, on the other hand it provides considerable evidence of their segmentation, of their being different from all other asset classes under consideration. The best R^2_{OOS} for a digital currency is 16% (for LTC), while the two other top R^2_{OOS} are all less than 9%; on the

Table 7. Recursive out-of-sample realized forecasting performances of alternative models.

		Recursive OLS Estimation		Best Recursive OLS Predictor			Dynamic Model Averaging		Dynamic Model Selection	
	Sample	OOS R2	RMSFE	Predictor	OOS R2	RMSFE	OOS R2	RMSFE	OOS R2	RMSFE
				Cryptocurrency returns						
Bitcoin returns	Jan. 2011 – Dec 2021	−0.5405	1.1869	Bitcoin Lagged Return	−0.0137	0.9628	−0.0875	0.997	**0.0748**	**0.9198**
Litecoin returns	July 2013 – Dec 2021	−0.2744	0.8901	Google searches	0.0111	0.7841	0.0974	0.748	**0.1619**	**0.7210**
Ripple returns	April 2015 – Dec 2021	−0.0439	1.0576	Google searches	0.0776	0.9941	0.0089	1.013	**0.0704**	**0.9965**
Ethereum returns	June 2016 – Dec 2021	−0.1565	1.1073	Google searches	0.0383	1.0097	−0.0318	1.033	**0.0894**	**0.9708**
				Benchmark asset returns						
Inv. Grade US Corporate bond returns	Jan. 2011 – Dec 2021	0.2068	0.8968	US Electrical Power Stocks	0.1794	0.9122	0.2238	0.8931	**0.3142**	**0.8394**
VW CRSP Mkt excess returns	Jan. 2011 – Dec 2021	0.3318	0.8238	US Electrical Power Stocks	0.3287	0.8251	0.4265	0.7664	**0.4842**	**0.7268**
VW World Developed Market returns (ex-US)	Jan. 2011 – Dec 2021	−0.2694	1.1328	Google_Bitcoin	−0.0088	1.0098	−0.0690	1.0387	**0.0674**	**0.9701**
Gold returns	Jan. 2011 – Dec 2021	−0.2420	1.1208	HML	−0.0172	1.0143	−0.0514	1.0332	**0.0927**	**0.9598**
VW US Dollar Exchange Rate Returns	Jan. 2011 – Dec 2021	0.2625	0.8610	Silver	0.2789	0.8514	0.3318	0.8215	**0.4362**	**0.7546**

Notes: The table reports two indicators of realized, OOS predictive accuracy: the root mean-squared forecast error and the OOS R-square. Four alternative forecasting models are compared, recursive OLS including all predictors jointly, (unfeasible) recursive OLS based on the best, full-sample predictor (that is also reported in the table), dynamic model averaging, and dynamic model selection. For each cryptocurrency/asset class, we have boldfaced the model yielding the lowest RMSFE and OOS R-square across the four predictive frameworks. The recursive, expanding window OOS experiments are applied to the available sample after initializing the recursive estimated on the basis of 52 initial observations.

opposite, all remaining asset classes (but international equity returns, ex-US) are characterized by R^2_{OOS}s that range between 9.3 (Gold returns) and a rather large 48 % (the excess returns on the aggregate stock market). Moreover, with no exception out of nine currencies/asset classes covered by Table 7, it is DMS that provides the most accurate OOS forecasting performance, probably as one would expect given the parsimonious nature of the algorithm; when this unfeasible model is taken out of the list, it is often the DMS to prevail.

One comment is in order. Many Readers may consider the values for the R^2_{OOS}s in the upper range of Table 7 as falling outside typical values in the predictability literature. Moreover, for instance in comparison to the classical paper by Goyal and Welch (2007), the R^2_{OOS}s yielded by simple recursive OLS appears to be relatively high. In Table C.1 in the Online Appendix, we show that such high R^2_{OOS}s essentially derive from the 2019–2020 sample. In fact, in Table B.1 the R^2_{OOS} for the traditional asset classes range between 0.3% and 6.2 %, while the scores yielded by recursive OLS are all deeply negative. Further analysis reveals that the recursive sample mean predictor used at the denominator of the definition of R^2_{OOS} massively under-performed in 2019–2020. In the case of cryptocurrencies, because of their declining prices during 2019. In the case of most traditional asset classes, because of the shock to prices brought forward by the Covid-19 pandemic shock. Implicitly, this supports the idea that the *relative* predictability of traditional asset classes is much stronger during crisis periods, such as the pandemic crisis, as emphasized, for instance, by Rapach, Strauss, and Zhou (2010b).

4.2. Realised OOS economic performance

As it is well known, assessing the predictability of asset returns under classical, statistical loss functions may turn out to be excessively remote when compared to the typical usage of these very forecasts in financial decisions, such as trading, portfolio allocation, and risk management, see e.g. Leitch and Tanner (1991). In fact, a number of papers (see, e.g. Cenesizoglu and Timmermann 2012 and Dal Pra et al. 2018) have emphasized that often the typical statistical loss functions used in much research on predictability may reveal weak forecasting power that is however able to generate substantial OOS, realized economic value under commonly used trading and asset allocation strategies. With *economic value*, we refer to the fact that enriched asset menus (for instance, to include alternative asset classes) or more sophistical asset allocation methods that exploit predictability ought to lead to superior OOS performances, for instance in terms of realized certainty equivalent returns, Sharpe ratios, and better higher-order properties of portfolio returns (skewness and kurtosis). The potential for a divergent informative content of statistical vs. economically grounded loss functions clearly represents an issue for our research design that has been so far based on statistical loss functions only.

As a result, in this section we proceed to extend our earlier evidence to a simple, representative portfolio problem to test whether exploiting the scant predictability in cryptocurrency returns may generate any economic value. In particular, we test two hypotheses. The first is that the inclusion of cryptocurrencies in the asset menu in addition to traditional assets classes ought to create substantial or, in any event, positive economic value even ignoring any predictability, just because of their segmentation from other assets. The second is that predictability can be exploited by an investor to achieve economic gains. Our portfolio allocation design is rather typical of the literature, see e.g. Barberis (2000) and Guidolin and Timmermann (2007). We model a US investor who, starting from unit wealth and in the presence of no-short sale constraints, maximizes expected power utility by selecting at weekly frequency the portfolio weights to be assigned to the N risky assets in the asset menu she faces:

$$\max_{\boldsymbol{\omega}_t} E_t^{\mathcal{M}}\left[\frac{W_{t+1}^{1-\gamma}}{1-\gamma}\right] \; (\gamma \neq 1) \quad \text{s.t. } W_{t+1} = \sum_{j=1}^{N} \omega_{j,t}(r_{j,t+1} - r_{f,t+1}) + (1 + r_{f,t+1}), \quad \omega_{j,t} \in [0,1]. \tag{7}$$

In problem (7), $r_{f,t+1}$ is the riskless, cash rate known at time t and the conditional expectation of the one-week ahead utility is computed by an IID bootstrap of the empirical distribution of the data up to time t, under a given prediction model $\mathcal{M}$.[13] For instance, under DMA, we shall have $W_{t+1}(\omega_t) \sim \mathcal{D}(\widehat{\mu}_{t+1\,|\,t}(\omega_t), \omega_t' Cov(\mathbf{r}_{t+1} - \hat{\boldsymbol{r}}^{DMA}_{t+1\,|\,t})\omega_t)$, where $\mathcal{D}$ is the empirical distribution of the available data obtained by a simple IID bootstrap from the available returns up to time t with mean $\widehat{\mu}_{t+1\,|\,t}(\omega_t) = \sum_{j=1}^{N} \omega_{j,t}(\hat{r}^{DMA}_{j,t+1\,|\,t} - r_{f,t+1}) + (1 + r_{f,t+1}) =$

$\sum_{j=1}^{N}\omega_{j,t}(\sum_{l=1}^{Q}\pi_{j,t+1\,|\,t,l}\hat{\boldsymbol{\theta}}_{j,t}'(l)\mathbf{z}_{j,t}-r_{f,t+1})+(1+r_{f,t+1})$ and variance $\omega_t' Cov(\mathbf{r}_{t+1}-\hat{\mathbf{r}}_{t+1\,|\,t}^{DMA})\omega_t$, where the latter is the $N\times N$ covariance matrix of the forecast errors from the DMA model.[14]

We recursively solve Equation (7) using standard constrained optimization methods over two alternative OOS periods: January 2012 – December 2021 and March 2020 – December 2021 (to capture the effects of the Covid-19 pandemic shock). We use one year of data to initialize the bootstrapping scheme; as far as the first OOS period is concerned, the initial bootstrap is based on observations from January 2011 to December 2011; for the second OOS period, the initial bootstrap is based on observations from March 2019 to February 2020. With reference to such OOS periods, we compute and report two measures of realized, risk-adjusted portfolio performance. The certainty equivalent return (CER) associated to the recursive solution of Equation (7) is computed as

$$CER(\mathcal{M},N)\equiv\left\{\frac{1}{T-P_N}\sum_{t=P_N+1}^{T}\left[\sum_{j=1}^{N}\hat{\omega}_{j,t}^{\mathcal{M}}(r_{j,t+1}-r_{f,t+1})+(1+r_{f,t+1})\right]^{1-\gamma}\right\}^{\frac{1}{1-\gamma}}-1, \tag{8}$$

where the OOS for the asset menu characterized by N assets starts at P_N+1 and the notation $CER(\mathcal{M},N)$ emphasizes that this the CER computed under the forecasts from model $\mathcal{M}$ for asset menu N. The Sharpe ratio is computed as

$$SR(\mathcal{M},N)\equiv\frac{\frac{1}{T-P_N}\sum_{t=P_N+1}^{T}\sum_{j=1}^{N}\hat{\omega}_{j,t}^{\mathcal{M}}(r_{j,t+1}-r_{f,t+1})}{\sqrt{\frac{1}{T-P_N}\sum_{t=P_N+1}^{T}\left[\sum_{j=1}^{N}\hat{\omega}_{j,t}^{\mathcal{M}}(r_{j,t+1}-r_{f,t+1})-\frac{1}{T-P_N}\sum_{t=P_N+1}^{T}\sum_{j=1}^{N}\hat{\omega}_{j,t}^{\mathcal{M}}(r_{j,t+1}-r_{f,t+1})\right]^2}}.$$

Of course, the standard caution that only the CER is really meaningful when portfolio weights have been optimized under Equation (7) also applies in this case, when asset returns are highly non-normal. However, reporting and commenting Sharpe ratios seems to be common practice, especially in the industry, and so we also follow this custom here. In practice, we consider three alternative values of the coefficient of constant relative risk aversion γ (i.e. 3, 7, and 15) and three alternative asset menus, which correspond to six distinct experiments: A first `baseline` menu characterized by $N_1=5$ risky assets, i.e. US stocks, World developed markets stocks (ex-US), US investment grade corporate bonds, gold, and long-short position in US dollars vis-á-vis a trade-weighted basket of other major currencies.[15] These are the five traditional asset classes that we have considered when performing comparisons with respect to the predictability patterns of cryptocurrencies. The portfolio calculations are performed under DMA, DMS, and when there is no predictability and returns are simply predicted by their historical sample means.

A second menu identical to the first, but expanded to include BTC, so that $N_2=6$; even though this second asset menu just includes one digital currency only, this is the most famous and by far the most actively traded (to the point of acting as a medium of exchange and currency conversion for other, minor crypto assets) that therefore allows us to apply our OOS portfolio tests the January 2012 – December 2021 sample. The portfolio calculations are performed under DMA, DMS, and when there is no predictability and returns are simply predicted by their historical sample means. A comparison of the CERs and Sharpe ratios of this asset menu vs. $N_1=5$ delivers an estimate of the economic value of adding Bitcoin to an otherwise traditional asset menu; a comparison of CERs and Sharpe ratios for $N_2=6$ provides instead an estimate of the economic value of capturing predictability when Bitcoin belongs to the asset menu.

A third asset menu, that further expands the second to include all cryptocurrencies, i.e. $N_3=9$; on the one hand, in the light of the goals of our paper, this expansion of the asset menu seems natural, also because there is a rising literature that has argued that often alternative cryptocurrencies display considerable cross-hedging power, as a reflection of moderate correlations, both unconditionally and conditional on a given coin being in a crash state (see, e.g. Baumöhl 2019); because of data availability, we only consider this asset menu for the shorter OOS period that covers the pandemic crisis, namely March 2020 – December 2021. A comparison of the CERs and Sharpe ratios of this asset menu vs. $N_1=5$ delivers an estimate of the economic value of adding all cryptocurrencies to an otherwise traditional asset menu; a comparison of CERs and Sharpe ratios for $N_3 = 9$ provides instead an estimate of the economic value of capturing predictability when all digital currencies jointly belong to the asset menu; finally a comparison of the CERs and Sharpe ratios of this asset menu vs. $N_2 = 6$ delivers an estimate of the economic value of adding Ethereum, Litecoin, and Ripple to an asset menu that includes only Bitcoin, which is another interesting question.

4.2.1. Portfolio weights

With reference to the case of $\gamma = 3$, Table 8 provides preliminary evidence on our dynamic portfolio exercises by reporting the customary set of summary statistics concerning optimal portfolio weights for the three asset menus described above. Tables D.1 and D.2 in the Online Appendix describe the empirical results for the cases of $\gamma = 7$ and 15 but are qualitatively similar, especially in the case of $\gamma = 7$, apart from the obvious finding that the optimal asset allocations in this cases are increasingly tilted away from the most volatile assets and towards gold and US dollars. The table is organized in two panels, devoted to results obtained under DMA (top) and DMS (bottom). For the sake of brevity, we do not report the portfolio allocation obtained by using simple historical, which remains available from the Authors. The main implications are still commented in the main text. When no cryptocurrencies are available, under DMA and DMS, on average 40% of the wealth is invested in US stocks, 5%–7% is invested in other developed stock markets and the rest is almost equally split across the other asset classes. However, all the weights display high volatilities (with a peak of 48% for the weights on the US equity) and there are occasional instances of the investor allocating all her wealth to just an asset class. When predictability is ignored, the asset allocation is aggressively tilted towards US stocks (the average weight is 81%). When bitcoin is added to the asset menu, the investor decreases her allocation to the US stocks to buy the cryptocurrency. In particular, under DMA and DMS, the investor allocates approximately 35% to the US stocks and 15% to the Bitcoin, while under no predictability the weights for US equity and the crypto assets are on average 25% and 68%, respectively. The remaining portfolio weights are largely unaffected. Finally, during the pandemic shock and when all cryptocurrencies are made available to our representative investor, we find that all allocations become somewhat tilted towards crypto assets (in particular ETH and XRP, under DMA, and Litecoin and Ripple under DMS). The allocation to equity drastically decreases from as much as 45% when crypto assets were not included in the asset menu to only 8%–14% when all the cryptocurrencies are available. Under historical moments, when predictability is ignored, the allocation becomes aggressively twisted towards BTC and ETH, with average weights totaling 60%. Gold receives a weight of 35% on average, while the rest of the assets barely enters the asset allocations. To a large extent, the empirical results show that the investor tends to substitute equity with cryptocurrencies when the latter become available both under predictability and without it. That said, the standard deviations of the optimal weights assigned to the digital currencies and the range of variation spanned by their maximum minus minimum estimates are large, with the latter spanning from 0 to 1. This indicates that such optimal shares are largely unstable over time and characterized by jumps and spikes, as it may be expected in the light of the patterns displayed by both the historical mean returns and the predictive coefficients in Tables 2–5.

4.2.2. Portfolio returns

Table 9 reports the portfolio realized statistics for $\gamma = 3, 7$, and 15 under the alternative asset menus with and without accounting for predictability. Notably, irrespective to the level of risk aversion of the investor, exploiting predictability seems to lead to better performances in terms of realized average returns, Sharpe ratios and certainty equivalent, in particular when DMS is applied. For instance, when the full sample period is considered, an investor with $\gamma = 3$ and access to Bitcoin would have obtained an average weekly return of 1.48% under DMS, which is more than twice what she would have achieved if she had ignored predictability. Despite portfolio returns are less volatile when predictability is neglected, an investor exploiting DMS forecasts would have obtained a Sharpe ratio of 0.21 (vs. 0.11 in the case of no predictability). Finally, under DMS the certainty equivalent return is 0.81%, while this only 0.18% under no predictability. This implies a value of predictability of 0.64% per week, which can be interpreted as the maximum fee that the investor is willing to pay to have access to the DMS forecasts.

Table 8. Summary statistics for recursive, expected power utility ($\gamma = 3$) portfolio weights.

	Mean	Median	St. Dev.	Min	Max	5th percentile	95th percentile
				Dynamic Model Average Portfolio Weights			
				Baseline (Full sample)			
CRSP VW Mkt	0.4325	0.0003	0.4810	0.0000	1.0000	0.0000	0.9999
Developed (ex-US) VW Mkt	0.0557	0.0000	0.1942	0.0000	1.0000	0.0000	0.5449
US IG Corporate Bonds	0.1513	0.0001	0.3431	0.0000	1.0000	0.0000	0.9999
Gold	0.1418	0.0000	0.3244	0.0000	1.0000	0.0000	0.9997
Long-short dollar trade	0.2186	0.0001	0.3979	0.0000	1.0000	0.0000	0.9999
				Baseline + Bitcoin (Full sample)			
Bitcoin	0.2059	0.0920	0.2683	0.0000	1.0000	0.0000	0.8232
CRSP VW Mkt	0.3429	0.0002	0.4140	0.0000	1.0000	0.0000	0.9999
Developed (ex-US) VW Mkt	0.0333	0.0000	0.1395	0.0000	0.9999	0.0000	0.2766
US IG Corporate Bonds	0.1292	0.0000	0.3033	0.0000	1.0000	0.0000	0.9796
Gold	0.1184	0.0000	0.2895	0.0000	1.0000	0.0000	0.9495
Long-short dollar trade	0.1703	0.0000	0.3398	0.0000	1.0000	0.0000	0.9996
				Baseline + All Crypto (Mar. 2020–Dec. 2021)			
Bitcoin	0.1007	0.0000	0.2313	0.0000	0.9999	0.0000	0.7599
Ethereum	0.2678	0.0003	0.3723	0.0000	1.0000	0.0000	0.9997
Litecoin	0.0814	0.0000	0.2425	0.0000	1.0000	0.0000	0.8223
Ripple	0.1767	0.0000	0.3518	0.0000	1.0000	0.0000	1.0000
CRSP VW Mkt	0.1405	0.0000	0.3017	0.0000	1.0000	0.0000	0.9792
Developed (ex-US) VW Mkt	0.0104	0.0000	0.1018	0.0000	0.9979	0.0000	0.0001
US IG Corporate Bonds	0.0042	0.0000	0.0269	0.0000	0.2127	0.0000	0.0002
Gold	0.1668	0.0000	0.3053	0.0000	1.0000	0.0000	0.9828
Long-short dollar trade	0.0515	0.0000	0.2058	0.0000	0.9999	0.0000	0.6188
				Dynamic Model Selection Portfolio Weights			
				Baseline (Full sample)			
CRSP VW Mkt	0.4154	0.0002	0.4814	0.0000	1.0000	0.0000	0.9999
Developed (ex-US) VW Mkt	0.0792	0.0000	0.2525	0.0000	1.0000	0.0000	0.9938
US IG Corporate Bonds	0.1609	0.0001	0.3543	0.0000	1.0000	0.0000	0.9998
Gold	0.1252	0.0000	0.3108	0.0000	1.0000	0.0000	0.9998
Long-short dollar trade	0.2193	0.0000	0.4056	0.0000	1.0000	0.0000	0.9999
				Baseline + Bitcoin (Full sample)			
Bitcoin	0.1643	0.0063	0.2591	0.0000	1.0000	0.0000	0.8334
CRSP VW Mkt	0.3418	0.0001	0.4252	0.0000	1.0000	0.0000	0.9999
Developed (ex-US) VW Mkt	0.0551	0.0000	0.1986	0.0000	0.9999	0.0000	0.5745
US IG Corporate Bonds	0.1419	0.0000	0.3245	0.0000	1.0000	0.0000	0.9995
Gold	0.1112	0.0000	0.2876	0.0000	1.0000	0.0000	0.9983
Long-short dollar trade	0.1856	0.0000	0.3667	0.0000	1.0000	0.0000	0.9998
				Baseline + All Crypto (Mar. 2020–Dec. 2021)			
Bitcoin	0.0043	0.0000	0.0420	0.0000	0.4114	0.0000	0.0000
Ethereum	0.0330	0.0000	0.1429	0.0000	0.9999	0.0000	0.2910
Litecoin	0.1426	0.0000	0.3033	0.0000	1.0000	0.0000	0.9845
Ripple	0.2066	0.0000	0.3566	0.0000	1.0000	0.0000	1.0000
CRSP VW Mkt	0.0849	0.0000	0.2623	0.0000	1.0000	0.0000	1.0000
Developed (ex-US) VW Mkt	0.4453	0.2796	0.4464	0.0000	1.0000	0.0000	1.0000
US IG Corporate Bonds	0.0104	0.0000	0.1021	0.0000	0.9999	0.0000	0.0000
Gold	0.0417	0.0000	0.2009	0.0000	1.0000	0.0000	0.0000
Long-short dollar trade	0.0313	0.0000	0.1749	0.0000	1.0000	0.0000	0.0000

Notes: The statistics are based on recursive, weekly optimized weights from a power utility with $\gamma = 3$. The percentile statistics are across time.

Other two considerations are in order. First, predictive models (and especially DMS) turn out to be valuable notwithstanding whether cryptocurrencies are included or not in the asset menu. For instance, an investor with $\gamma = 3$, without access to Bitcoin, would still obtain a gain of 0.62% in terms of certainty equivalent return by exploiting DMS forecasts. Second, the economic gain from predictability seems to be larger during the Covid-19 pandemic. For instance, with $gamma = 3$, exploiting DMS forecasts during the Covid-19 period will increase the CER by 2.7% when cryptocurrencies are available. This is in line with the literature (see, e.g. Rapach, Strauss,

Table 9. Realized performance of recursive, expected power utility portfolio weights for three alternative values of the coefficient of relative risk aversion and three alternative asset menus/sample sizes, the table reports realized performance measures for optimal portfolio weights derived from the maximization of expected power utility.

	Mean	Volatility	Sharpe ratio	Skewness	Kurtosis	Realized avg. utility	CER	Value of Predict.	Value of crypto
					$\gamma = 3$				
					Baseline (Full sample)				
DMA	0.0067	0.0159	0.4125	1.4961	**13.4019**	−0.4938	0.0063		
DMS	**0.0081**	**0.0157**	**0.5083**	**2.0160**	14.4142	−**0.4924**	**0.0077**		
Historical Mean	0.0020	0.0178	0.1087	−1.4367	26.1954	−0.4985	0.0015	0.0062	
					Baseline + Bitcoin (Full sample)				
DMA	0.0102	0.0695	0.1454	3.4274	37.8915	−0.4961	0.0039		
DMS	**0.0148**	0.0710	**0.2064**	**3.6067**	39.9537	−**0.4920**	**0.0081**		
Historical Mean	0.0064	**0.0543**	0.1153	0.5697	**18.4081**	−0.4982	0.0018	0.0064	0.0004
					Baseline (Mar. 2020 – Dec. 2021)				
DMA	**0.0123**	0.0255	**0.4772**	1.4104	8.2250	−**0.4888**	**0.0114**		
DMS	0.0091	0.0256	0.3519	**1.4123**	8.3353	−0.4919	0.0082		
Historical Mean	0.0008	**0.0233**	0.0313	−0.8060	**6.0328**	−0.5000	0.0000	0.0113	
					Baseline + All Crypto (Mar. 2020 – Dec. 2021)				
DMA	0.0507	0.1832	0.2761	3.0211	17.5533	−0.4825	0.0180		
DMS	**0.0613**	0.1806	**0.3391**	**3.3370**	17.9020	−**0.4688**	**0.0327**		
Historical Mean	0.0195	**0.0913**	0.2123	−0.2753	**6.4015**	−0.4942	0.0059	0.0268	0.0245
					$\gamma = 7$				
					Baseline (Full sample)				
DMA	0.0067	0.0150	0.4398	1.8427	15.7852	−0.1608	0.0060		
DMS	**0.0082**	**0.0150**	**0.5371**	**2.3230**	16.7339	−**0.1594**	**0.0074**		
Historical Mean	0.0013	0.0117	0.0992	−1.9531	**13.9994**	−0.1659	0.0007	0.0067	
					Baseline + Bitcoin (Full sample)				
DMA	0.0080	0.0374	0.2098	0.2903	36.9329	−0.1660	0.0007		
DMS	**0.0114**	0.0405	**0.2786**	**2.3929**	37.5808	−**0.1615**	**0.0052**		
Historical Mean	0.0033	**0.0281**	0.1153	−0.5776	**25.7943**	−0.1667	0.0000	0.0052	−0.0022
					Baseline (Mar. 2020 – Dec. 2021)				
DMA	**0.0115**	0.0238	**0.4804**	**1.5973**	9.6183	−**0.1572**	**0.0097**		
DMS	0.0090	0.0249	0.3593	1.3606	7.9671	−0.1598	0.0071		
Historical Mean	0.0004	**0.0186**	0.0185	−1.1938	**6.9973**	−0.1675	−0.0008	0.0106	
					Baseline + All Crypto (Mar. 2020 – Dec. 2021)				
DMA	0.0464	0.1443	0.3208	3.5863	20.4881	−0.1561	0.0110		
DMS	**0.0545**	0.1597	**0.3405**	**4.4662**	27.9382	−**0.1469**	**0.0212**		
Historical Mean	0.0077	**0.0456**	0.1660	−0.4610	**5.2097**	−0.1668	−0.0002	0.0214	0.0142
					$\gamma = 15$				
					Baseline (Full sample)				
DMA	0.0065	0.0132	0.4817	2.7521	21.8289	−0.0663	0.0054		
DMS	**0.0078**	0.0137	**0.5646**	**3.1844**	22.3222	−**0.0651**	**0.0067**		
Historical Mean	0.0010	**0.0075**	0.1161	−1.9168	**13.9416**	−0.0709	0.0005	0.0062	
					Baseline + Bitcoin (Full sample)				
DMA	0.0071	0.0228	0.3076	−0.9638	37.8608	−0.0732	−0.0017		
DMS	**0.0096**	0.0230	**0.4114**	**2.1395**	24.0100	−**0.0659**	**0.0058**		
Historical Mean	0.0020	**0.0141**	0.1329	−0.5235	**20.7412**	−0.0712	0.0003	0.0055	−0.0009
					Baseline (Mar. 2020 – Dec. 2021)				
DMA	**0.0099**	0.0220	**0.4454**	**2.1337**	12.6823	−**0.0648**	**0.0069**		
DMS	0.0077	0.0180	0.4237	0.6491	**5.4583**	−0.0662	0.0054		
Historical Mean	0.0003	**0.0132**	0.0119	−1.3769	8.8091	−0.0726	−0.0012	0.0081	
					Baseline + All Crypto (Mar. 2020 – Dec. 2021)				
DMA	0.0314	0.0830	**0.3772**	3.1995	15.9061	−0.0643	0.0076		
DMS	**0.0361**	0.1037	0.3476	**3.9708**	20.5761	−**0.0615**	**0.0107**		
Historical Mean	0.0032	**0.0230**	0.1351	−0.7736	**5.2297**	−0.0727	−0.0012	0.0120	0.0053

Notes: All measures are reported on a weekly basis.

and Zhou 2010b) that has found larger evidence of predictability during economic downturns. The results are qualitatively similar for $\gamma = 7$ and $\gamma = 15$.

Table 9 also allows to answer the question that is central to this paper, i.e. what is the value of a seemingly segmented asset–digital currencies, at least in a predictive framework–in an asset management perspective. By taking the difference between the highest CER when crypto assets are included in the asset menu and the highest CER under the baseline menu, we find an increase modest in CER of 0.04% per week, when $\gamma = 3$ and a decrease in CER of 0.02% and 0.09% per week for $\gamma = 7$ and $\gamma = 15$, respectively. However, during the Covid-pandemic, a risk-averse investor would have benefited from the access to the four cryptocurrencies. In particular, in this sample period, access to cryptocurrencies would have increased the CER of an investor with $gamma = 3$ by 2.45% under the best predictability model. This points toward the fact that cryptocurrencies may have acted as a hedge during the Covid crisis. However, the benefit from the inclusion of cryptocurrencies decreases with risk aversion (the increase in CER is equal to 1.42% and 0.53% for $gamma = 7$ and $gamma = 15$, respectively).

Finally, by comparing the CERs of the baseline menu and the menu with crypto assets, we notice that an investor would benefit much more from the inclusion of cryptocurrencies when predictability is accounted for. For instance, during the Covid period, under no predictability an investor with $\gamma = 3$ would have obtained a gain in terms of CER equal to 0.59% when including cryptocurrencies in her portfolio. This is considerably smaller than the gain of 2.45% that is achieved when DMS forecasts are exploited.

5. Conclusion

In this paper, we have investigated the characteristics of cryptocurrencies taking a forecasting perspective.[16] Specifically, we have adopted a flexible, dynamic forecasting approach in which – given a set of plausible predictive variables drawn from earlier literature on the asset pricing of cryptocurrencies – we ask whether the patterns, the strength, and the economic value of any predictability characterizing cryptocurrency returns may differ from that typical of traditional asset classes.[17] Our approach is flexible because instead of simply performing recursive OLS estimation of linear predictive models with fixed predictors, we allow the data to either recursively re-weight the forecasting variables to used in the model or to select which variables ought to be dynamically included in the model (see Raftery, Kárnỳ, and Ettler 2010 and Koop and Korobilis 2012). Besides, we measure economic value through standard recursive asset allocation exercises in which a US investor maximizes expected power utility across alternative asset menus – with and without cryptocurrencies – and predictability models, including a sample historical mean benchmark that does not feature any predictability of asset returns. This exercise ensures adequate robustness to the selection of specific statistical loss functions, which remains a tricky choice in all research designs based on relative forecasting power.

We find evidence that digital currencies may indeed represent a new asset class, substantially segmented from traditional asset classes. Cryptocurrency appear to be: (1) characterized by returns that are generally less predictable on average when compared to other asset classes, including gold and the external, trade value-weighted US dollar; (2) characterized by returns that can be forecast according to patterns and with a measurable degree of time variation that differ from most other asset classes, including gold that has been often indicated as the most closely related asset class; (3) able to generate some modest, realized OOS economic value (especially when measured in terms of ex-post Sharpe ratios) when they are added to otherwise traditional asset menus of cash, corporate bonds, US and international stocks, and long-short exchange rate positions; (4) such economic benefit tend to be higher during the Covid-19 crisis. As always, an operational understanding of what is driving the segmentation would be more satisfying on logical grounds, and also allow us to forecast the long-term evolution of the segmentation we have uncovered. We leave this interesting question for future research.

Notes

1. Standard asset pricing theory states that, in the absence of arbitrage opportunities and of segmentation, the same pricing kernel(s) will price all asset classes, see, e.g. Schilling and Uhlig (2019). In its turn, this implies that all predictors that can forecast the pricing kernel ought to forecast excess returns on all assets.
2. The issues caused by structural change in the risk-return trade-off characterizing cryptocurrency returns has also been investigated by Krueckeberg and Scholz (2019).

3. For instance, in the case in which the models are defined according to whether each predictor is included or excluded, $Q = 2^P$ so that Π will be a $2^P \times 2^P$ matrix and will contain (because of sum up constraints) 2^{2P-1} parameters to estimate. For instance, with $P = 17$, we face $2^{33} = 8,589,934,592$ free parameters to estimate and the problem becomes computationally unfeasible.
4. To qualify for inclusion in the index, securities must have an investment grade rating (based on an average of Moody's, S&P, and Fitch) and an investment grade rated currency of denomination (based on an average of Moody's, S&P, and Fitch foreign currency-denominated long term sovereign debt ratings). Each security must have a residual maturity in excess of one year, a fixed coupon schedule, and a minimum amount outstanding of $250 million.
5. Mining is the process of adding transaction records to a cryptocurrency public ledger of past transactions. The primary purpose of mining is to set the history of transactions in a way that makes it computationally impractical to be altered by any one entity, preventing frauds. Mining is intentionally designed to be resource-intensive and difficult so that the number of blocks found each day by miners remains under control. In fact, limits to the supply are endogenously set. Mining is also the mechanism used to introduce new crypto coins into the system: miners are paid transaction fees as well as a 'subsidy' for each newly created coins.
6. As a robustness check, we have performed a selection of our analyses using the rate of growth in the Cambridge Bitcoin Energy Consumption Index, which takes into account the cost of producing Bitcoin vis-á-vis the efficiency of mining equipment and the costs of running mining facilities. We find similar results.
7. A constant intercept is always included in all models.
8. Interestingly, we found many instances in which the probability of inclusion may be non-negligible and yet the corresponding DMA predictive coefficients are modest in absolute value. This is possible when only a few models that include a given predictor with non-extreme coefficients receive a relatively high probability while the majority of models estimate a zero coefficient associated to the same predictor.
9. A recent literature has established that in a OOS forecasting perspective, recursive OLS are outperformed by weighting regression estimates by a class of time-dependent functions that assign heavier weights to more recent observations. DMA possesses such features, besides applying shrinking.
10. The set of predicted model/state probabilities under DMA and DMS are identical by construction because DMS is a probability-based refinement of DMA. Therefore the same analysis of Figure 1 performed above applies here. Figure B.2 in the Online Appendix reports a few plots of recursively estimated coefficients under DMS.
11. The full set of parameters estimates is available upon request from the Authors.
12. In unreported plots, we obtain ex-post DMA evidence of high (exceeding 0.6) and time-varying probability of inclusion for the SMB, HML, stock momentum, platinum, silver, and especially aggregate market returns. Interestingly, past bitcoin momentum predicts the effective US dollar exchange rate.
13. Because we impose that the empirical distribution of the data, means, variance and covariances are affected by the selected model $\mathcal{M}$, this IID bootstrap scheme to optimize portfolio weights is also known as a filtered historical simulation approach. In our view, in the presence of the massive deviation from normality characterizing the returns on cryptocurrencies, any other parametric approach to optimal asset allocation would be of dubious applicability.
14. Because we are considering the conditional, one-step ahead distribution of wealth, this is the relevant definition of covariance for our problem, see Campbell et al. (2004).
15. Because we assume that the short position is 100% covered by depositing cash, we impute on the long-short exchange rate the same cost in terms of commitment of available wealth as other assets/strategies.
16. Of course, if we were to set $\psi = 1$, then $\pi_{t+1\,|\,t,q}$ would be simply proportional to the marginal likelihood using data through time t, which is a standard Bayesian model averaging (BMA) approach. However, one formally has BMA only when $\lambda = \psi = 1$ which implies using conventional linear forecasting models with no time variation in coefficients.
17. Note that $\Sigma_{1|1}(q) = \Sigma_{1\,|\,0}(q) - \Sigma_{1\,|\,0}(q)\mathbf{z}_0(h_1(q) + \mathbf{z}_0'\Sigma_{1\,|\,0}(q)\mathbf{z}_0)^{-1}\mathbf{z}_0'\Sigma_{1\,|\,0}(q)$ just requires knowledge of $\Sigma_{1\,|\,0}(q) = \frac{1}{\lambda}\Sigma_{0\,|\,0}(q)$. We set $\Sigma_{0\,|\,0}(q) = 100\hat{h}_0(q)\mathbf{I}_{P+1}$, where the last estimate is specified in the main text. This prior is of course, very diffuse and reflects considerable, initial uncertainty on the financial nature of crytpocurrencies as an asset class.

Acknowledgments

We thank Dirk Baur (a discussant), Victor De Miguel, Gary Koop, Anestis Ladas, Emmanouil Platanakis, seminar partecipants at the University of Bath AFL – QF and Fintech Research Seminar Series, conference participants at the First Annual Conference of the International Association of Applied Econometrics, the 13th Conference on Computational and Financial Econometrics in 2019 (London), to CEMA 2021, and to the 13th Virtual Annual Conference of the Society for Financial Econometrics. We would also like to thank Volodymyr Ferenchenko and Federico Mainardi for valuable research assistance. A previous version of the paper was circulating under the title 'Dissecting Time-Varying Risk Exposures in Cryptocurrency Markets'.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Notes on contributors

Daniele Bianchi is an Associate Professor of Finance at the School of Economics and Finance, Queen Mary, University of London.

Massimo Guidolin is a Full Professor of Finance in the Department of Finance at Bocconi University.

Manuela Pedio is an Assistant Professor of Finance at the School of Accounting and Finance, University of Bristol.

ORCID

Daniele Bianchi http://orcid.org/0000-0002-2321-2108

References

Atsalakis, G. S., I. G. Atsalaki, F. Pasiouras, and C. Zopounidis. 2019. "Bitcoin Price Forecasting with Neuro-Fuzzy Techniques." *European Journal of Operational Research* 276 (2): 770–780.

Barberis, N. 2000. "Investing for the Long Run When Returns are Predictable." *The Journal of Finance*55 (1): 225–264.

Baumöhl, E. 2019. "Are Cryptocurrencies Connected to Forex? A Quantile Cross-spectral Approach." *Finance Research Letters* 29: 363–372.

Baur, D. G., K. Hong, and A. D. Lee. 2018. "Bitcoin: Medium of Exchange or Speculative Assets?" *Journal of International Financial Markets, Institutions and Money* 54: 177–189.

Bianchi, D., and M. Babiak. 2022. "A Factor Model for Cryptocurrency Returns." Available at SSRN 3935934.

Bianchi, D., M. Babiak, and A. Dickerson. 2022. "Trading Volume and Liquidity Provision in Cryptocurrency Markets." Working Paper.

Bianchi, D., M. Guidolin, and F. Ravazzolo. 2017. "Macroeconomic Factors Strike Back: A Bayesian Change-Point Model of Time-Varying Risk Exposures and Premia in the Us Cross-section." *Journal of Business & Economic Statistics* 35 (1): 110–129.

Binsbergen, V., H. Jules, and R. S. Koijen. 2010. "Predictive Regressions: A Present-Value Approach." *The Journal of Finance* 65 (4): 1439–1471.

Bouri, E., R. Gupta, A. Lahiani, and M. Shahbaz. 2018a. "Testing for Asymmetric Nonlinear Short- and Long-Run Relationships Between Bitcoin, Aggregate Commodity and Gold Prices." *Resources Policy*57: 224–235.

Bouri, E., R. Gupta, C. K. M. Lau, D. Roubaud, and S. Wang. 2018b. "Bitcoin and Global Financial Stress: A Copula-Based Approach to Dependence and Causality in the Quantiles." *The Quarterly Review of Economics and Finance* 69: 297–307.

Bouri, E., C. K. M. Lau, B. Lucey, and D. Roubaud. 2019. "Trading Volume and the Predictability of Return and Volatility in the Cryptocurrency Market." *Finance Research Letters* 29: 340–346.

Campbell, J. Y., G. Chacko, J. Rodriguez, and L. M. Viceira. 2004. "Strategic Asset Allocation in a Continuous-Time Var Model." *Journal of Economic Dynamics and Control* 28 (11): 2195–2214.

Cenesizoglu, T., and A. Timmermann. 2012. "Do Return Prediction Models Add Economic Value?" *Journal of Banking & Finance* 36 (11): 2974–2987.

Cheah, J. E. -T., D. Luo, Z. Zhang, and M. -C. Sung. 2022. "Predictability of Bitcoin Returns." *The European Journal of Finance* 28 (1): 66–85.

Chuen, K., L. David, L. Guo, and Y. Wang. 2017. "Cryptocurrency: A New Investment Opportunity?" Working Paper.

Corbet, S., A. Meegan, C. Larkin, B. Lucey, and L. Yarovaya. 2018. "Exploring the Dynamic Relationships Between Cryptocurrencies and Other Financial Assets." *Economics Letters* 165: 28–34.

Da, Z., J. Engelberg, and P. Gao. 2011. "In Search of Attention." *The Journal of Finance* 66 (5): 1461–1499.

Dal Pra, G., M. Guidolin, M. Pedio, and F. Vasile. 2018. "Regime Shifts in Excess Stock Return Predictability: An Out-of-Sample Portfolio Analysis." *The Journal of Portfolio Management* 44 (3): 10–24.

Dangl, T., and M. Halling. 2012. "Predictive Regressions with Time-Varying Coefficients." *Journal of Financial Economics* 106 (1): 157–181.

Detzel, A., H. Liu, J. Strauss, G. Zhou, and Y. Zhu. 2021. "Learning and Predictability Via Technical Analysis: Evidence from Bitcoin and Stocks with Hard-to-Value Fundamentals." *Financial Management*50 (1): 107–137.

Drobetz, W., P. P. Momtaz, and H. Schröder. 2019. "Investor Sentiment and Initial Coin Offerings." *The Journal of Alternative Investments* 21 (4): 41–55.

Dyhrberg, A. H. 2016a. "Bitcoin, Gold and the Dollar–a Garch Volatility Analysis." *Finance Research Letters* 16: 85–92.

Dyhrberg, A. H. 2016b. "Hedging Capabilities of Bitcoin. Is it the Virtual Gold?" *Finance Research Letters* 16: 139–144.

Fama, E. F., and K. R. French. 2015. "A Five-Factor Asset Pricing Model." *Journal of Financial Economics* 116 (1): 1–22.

Goyal, A., and I. Welch. 2007. "A Comprehensive Look at the Empirical Performance of Equity Premium Prediction." *The Review of Financial Studies* 21 (4): 1455–1508.

Guidolin, M. 2011. "Markov Switching Models in Empirical Finance." *Advances in Econometrics* 27 (2): 1–86.

Guidolin, M., E. Hansen, and M. Pedio. 2019. "Cross-asset Contagion in the Financial Crisis: A Bayesian Time-varying Parameter Approach." *Journal of Financial Markets* 45: 83–114.

Guidolin, M., and A. Timmermann. 2007. “Asset Allocation under Multivariate Regime Switching.” *Journal of Economic Dynamics and Control* 31 (11): 3503–3544.
Hudson, R., and A. Urquhart. 2021. “Technical Trading and Cryptocurrencies.” *Annals of Operations Research* 297 (1): 191–220.
Johannes, M., A. Korteweg, and N. Polson. 2014. “Sequential Learning, Predictability, and Optimal Portfolio Returns.” *The Journal of Finance* 69 (2): 611–644.
Kajtazi, A., and A. Moro. 2019. “The Role of Bitcoin in Well Diversified Portfolios: A Comparative Global Study.” *International Review of Financial Analysis* 61: 143–157.
Klein, T., H. P. Thu, and T. Walther. 2018. “Bitcoin is Not the New Gold–a Comparison of Volatility, Correlation, and Portfolio Performance.” *International Review of Financial Analysis* 59: 105–116.
Koop, G., and D. Korobilis. 2012. “Forecasting Inflation Using Dynamic Model Averaging.” *International Economic Review* 53 (3): 867–886.
Krueckeberg, S., and P. Scholz. 2019. “Cryptocurrencies as an Asset Class.” In *Cryptofinance and Mechanisms of Exchange*, 1–28.
Leitch, G., and J. E. Tanner. 1991. “Economic Forecast Evaluation: Profits Versus the Conventional Error Measures.” *The American Economic Review* 81: 580–590.
Li, J., N. Li, J. Peng, H. Cui, and Z. Wu. 2019. “Energy Consumption of Cryptocurrency Mining: A Study of Electricity Consumption in Mining Cryptocurrencies.” *Energy* 168: 160–168.
Li, J., and G. Yi. 2019. “Toward a Factor Structure in Crypto Asset Returns.” *The Journal of Alternative Investments* 21 (4): 56–66.
Liu, Y., and A. Tsyvinski. 2020. “Risks and Returns of Cryptocurrency.” *The Review of Financial Studies*34 (6): 2689–2727.
Pastor, L., and R. F. Stambaugh. 2009. “Predictive Systems: Living with Imperfect Predictors.” *The Journal of Finance* 64 (4): 1583–1628.
Pesaran, M. H., and A. Timmermann. 2005. “Small Sample Properties of Forecasts from Autoregressive Models under Structural Breaks.” *Journal of Econometrics* 129 (1-2): 183–217.
Pettenuzzo, D., A. Timmermann, and R. Valkanov. 2014. “Forecasting Stock Returns under Economic Constraints.” *Journal of Financial Economics* 114 (3): 517–553.
Plakandaras, V., E. Bouri, and R. Gupta. 2021. “Forecasting Bitcoin Returns: Is There a Role for the US–China Trade War?” *Journal of Risk* 23 (3).
Raftery, A. E., M. Kárnỳ, and P. Ettler. 2010. “Online Prediction under Model Uncertainty Via Dynamic Model Averaging: Application to a Cold Rolling Mill.” *Technometrics* 52 (1): 52–66.
Rapach, D., J. Strauss, and G. Zhou. 2010a. “Out-of-Sample Equity Premium Prediction: Combination Forecasts and Links to the Real Economy.” *Review of Financial Studies* 23 (2): 821–862.
Rapach, D. E., J. K. Strauss, and G. Zhou. 2010b. “Out-of-Sample Equity Premium Prediction: Combination Forecasts and Links to the Real Economy.” *The Review of Financial Studies* 23 (2): 821–862.
Schilling, L., and H. Uhlig. 2019. “Some Simple Bitcoin Economics.” *Journal of Monetary Economics*106: 16–26.
Stock, J. H., and M. W. Watson. 1996. “Evidence on Structural Instability in Macroeconomic Time Series Relations.” *Journal of Business & Economic Statistics* 14 (1): 11–30.
Urom, C., I. Abid, K. Guesmi, and J. Chevallier. 2020. “Quantile Spillovers and Dependence Between Bitcoin, Equities and Strategic Commodities.” *Economic Modelling* 93: 230–258.
Yermack, D. 2015. “Is Bitcoin a Real Currency? An Economic Appraisal.” In *Handbook of Digital Currency*, 31–43.

Appendix. Estimation strategy

For given values of h_t and $\mathbf{L}_t$, standard Kalman filtering starts with $\theta_t|r_t, r_{t-1}, \ldots, r_1 \sim N(\hat{\boldsymbol{\theta}}_t, \Sigma_{t|t})$, where $\Sigma_{t|t}$ depends on h_t and $\mathbf{Q}_t$. Then filtering proceeds using:

$$\boldsymbol{\theta}_{t+1}|r_t, r_{t-1}, \ldots, r_1 \sim N(\hat{\boldsymbol{\theta}}_t, \boldsymbol{\Sigma}_{t+1\,|\,t}) \quad \boldsymbol{\Sigma}_{t+1\,|\,t} = \boldsymbol{\Sigma}_{t|t} + \mathbf{L}_{t+1}. \tag{A1}$$

Raftery, Kárnỳ, and Ettler (2010) have noted that the computational burden simplifies substantially when the updating recursion is simplified to:

$$\boldsymbol{\Sigma}_{t+1\,|\,t} = \frac{1}{\lambda}\boldsymbol{\Sigma}_{t\,|\,t} = \boldsymbol{\Sigma}_{t\,|\,t} + \left(\frac{1}{\lambda} - 1\right)\boldsymbol{\Sigma}_{t\,|\,t}, \tag{A2}$$

or equivalently $\mathbf{L}_{t+1} = (\frac{1}{\lambda}) - 1)\Sigma_{t\,|\,t}$ where $0 < \lambda \leq 1$. Importantly, this formula does not depend on the estimate of $\mathbf{L}_{t+1}$. In line with the empirical literature, we set $\lambda = 0.99$. Estimation in the one model case is then completed by the updating equations:

$$\boldsymbol{\theta}_{t+1} \,|\, r_{t+1}, r_t, r_{t-1}, \ldots, r_1 \sim N(\hat{\boldsymbol{\theta}}_{t+1}, \boldsymbol{\Sigma}_{t+1\,|\,t+1}) \tag{A3}$$

$$\hat{\boldsymbol{\theta}}_{t+1} = \hat{\boldsymbol{\theta}}_t + \boldsymbol{\Sigma}_{t+1\,|\,t}\mathbf{z}_t(h_{t+1} + \mathbf{z}_t'\boldsymbol{\Sigma}_{t+1\,|\,t+1}\mathbf{z}_t)^{-1}(r_{t+1} - {\boldsymbol{\theta}_t}'\mathbf{z}_t) \tag{A4}$$

$$\boldsymbol{\Sigma}_{t+1\,|\,t+1} = \boldsymbol{\Sigma}_{t+1\,|\,t} - \boldsymbol{\Sigma}_{t+1\,|\,t}\mathbf{z}_t(h_{t+1} + \mathbf{z}_t'\boldsymbol{\Sigma}_{t+1\,|\,t}\mathbf{z}_t)^{-1}\mathbf{z}_t'\boldsymbol{\Sigma}_{t+1\,|\,t} \tag{A5}$$

$$r_{t+1}|r_t, r_{t-1}, \ldots, r_1 \sim N(\hat{\boldsymbol{\theta}}_t'\mathbf{z}_t, h_t + \mathbf{z}_t'\boldsymbol{\Sigma}_{t+1\,|\,t}\mathbf{z}_t). \tag{A6}$$

Conditional on h_t, these results are analytical and, thus, no Markov chain Monte Carlo (MCMC) algorithm is required. This greatly reduces the computational burden. In the case of DMA, we use this approximation of the Kalman filter and an additional one. Call Θ_t the vector collecting all coefficients, $\Theta_t \equiv [\theta_t'(1)\ \theta_t'(2) \cdots \theta_t'(Q)]'$. First, the previous filtering equations for θ_{t+1} now become:

$$\Theta_t \mid M_t = q, r_t, r_{t-1}, \ldots, r_1 \sim N(\hat{\boldsymbol{\theta}}_t(q), \boldsymbol{\Sigma}_{t\,|\,t}(q)) \tag{A7}$$

$$\Theta_{t+1} \mid M_{t+1} = q, r_t, r_{t-1}, \ldots, r_1 \sim N(\hat{\boldsymbol{\theta}}_t(q), \boldsymbol{\Sigma}_{t+1\,|\,t}(q)) \tag{A8}$$

$$\Theta_{t+1} \mid M_{t+1} = q, r_{t+1}, r_t, r_{t-1}, \ldots, r_1 \sim N(\hat{\boldsymbol{\theta}}_{t+1}(q), \boldsymbol{\Sigma}_{t+1\,|\,t+1}(q)), \tag{A9}$$

where $\hat{\boldsymbol{\theta}}_t(q)$, $\Sigma_{t|t}(q)$, and $\Sigma_{t+1\,|\,t}(q)$ are computed from (A4), (A5), and (A2), respectively. However, conditional on $M_t = q$, the prediction and updating equations will only provide information on $\theta_t(q)$ and not the full vector Θ_t: one needs a method for unconditional prediction to obtain probability-weighting of forecasts across values for $\hat{\boldsymbol{\theta}}_t(q)$. Koop and Korobilis (2012) introduce one additional, forgetting-like factor to be applied to the state equation across models, ψ, comparable to the forgetting factor used with the state equation for the parameters, λ. Note that by definition the probability transition equation is,

$$\begin{aligned} pdf(\Theta_t \mid r_t, r_{t-1}, \ldots, r_1) &= \sum_{q=1}^{Q} pdf(\boldsymbol{\theta}_t \mid M_t = q, r_t, r_{t-1}, \ldots, r_1)\, \Pr(M_t = q, r_t, r_{t-1}, \ldots, r_1) \\ &= \sum_{q=1}^{Q} pdf(\boldsymbol{\theta}_t \mid M_t = q, r_t, r_{t-1}, \ldots, r_1) \pi_{t|t,q}, \end{aligned} \tag{A10}$$

where $pdf(\theta_t \mid M_t = q, r_t, r_{t-1}, \ldots, r_1)$ comes from $\theta_t(q) \mid r_t, r_{t-1}, \ldots, r_1 \sim N(\hat{\boldsymbol{\theta}}_t(q), \Sigma_{t\,|\,t}(q))$. In the case of a standard Markov switching model, when the elements of the transition matrix $\mathbf{T}$ are $\tau_{ij} = \Pr(M_{t+1} = j \mid M_t = i)$, (A10) would be:

$$\pi_{t+1\,|\,t,q} = \sum_{l=1}^{Q} \pi_{t\,|\,t,l} \tau_{lq}, \tag{A11}$$

but we approximate it by:

$$\pi_{t+1\,|\,t,q} = \frac{\pi_{t\,|\,t,q}^{\psi}}{\sum_{l=1}^{Q} \pi_{t\,|\,t,l}^{\psi}} \quad 0 < \psi \leq 1 \tag{A12}$$

which is a type of multi-parameter power steady model. To understand further how the forgetting factor can be interpreted, note that this specification implies that the weight used in DMA which is attached to model q at time t is updated according to:

$$\begin{aligned} \pi_{t+1\,|\,t,q} &= \frac{[\pi_{t\,|\,t-1,q} pdf(r_{t-1} \mid M_{t-1} = q, r_{t-2}, \cdots, r_1)]^{\psi}}{\sum_{l=1}^{Q} [\pi_{t\,|\,t-1,l} pdf(r_{t-1} \mid M_{t-1} = l, r_{t-2}, \cdots, r_1)]^{\psi}} \\ &\propto [\pi_{t-1\,|\,t-2,q} pdf(r_{t-2} \mid M_{t-2} = q, r_{t-3}, \cdots, r_1)]^{\psi^2} [pdf(r_{t-1} \mid M_{t-1} = q, r_{t-2}, \cdots, r_1)]^{\psi} \\ &\propto \cdots \propto \prod_{i=1}^{t-1} [pdf(r_{t-i} | M_{t-i} = q, r_{t-i-1}, \cdots, r_1)]^{\psi^i} \pi_{0\,|\,0,q}^{\psi^t}, \end{aligned} \tag{A13}$$

so that model q receive more weight at time t if it has forecast well in the recent past (where forecast performance is measured by the predictive density, $pdf(r_{t-i} \mid M_{t-i} = q, r_{t-i-1}, \cdots, r_1)$). The interpretation of 'recent past' is controlled by the forgetting factor, and we have the same exponential decay at the rate ψ^i for observations i periods ago. Clearly, it is not enough for one time performance in the past to highly disappointing, i.e. $pdf(r_{t-i} \mid M_{t-i} = q, r_{t-i-1}, \cdots, r_1) \simeq 0$, for $\pi_{t+1\,|\,t,q} \simeq 0$ because the expression in (A13) fails to iterate over a scaling factor which may also get very close to zero when in the past all the moments have produced some performances characterized by very small predictive density scores.[16] In line with the empirical literature, and also for consistency with our earlier choice for λ, we set $\psi = 0.99$.

The benefit of using this approximation in the model prediction equation is that we do not require an MCMC algorithm to draw transitions between models nor a simulation algorithm over the space of models. The reason is that a simple, but effective updating equation is:

$$\pi_{t\,|\,t,q} = \frac{\pi_{t\,|\,t-1,q} \cdot \overbrace{pdf(r_t \mid M_t = q, r_{t-1}, \ldots, r_1)}^{\text{from} \sim N(\hat{\boldsymbol{\theta}}_{t-1}'(q)\mathbf{z}_{t-1}, h_t + \mathbf{z}_{t-1}' \boldsymbol{\Sigma}_{t\,|\,t-1}(q)\mathbf{z}_{t-1})}}{\sum_{l=1}^{Q} \pi_{t\,|\,t-1,l} \cdot \underbrace{pdf(r_t \mid M_t = l, r_{t-1}, \ldots, r_1)}_{\text{from} \sim N(\hat{\boldsymbol{\theta}}_{t-1}'(l)\mathbf{z}_{t-1}, h_t + \mathbf{z}_{t-1}' \boldsymbol{\Sigma}_{t\,|\,t-1}(l)\mathbf{z}_{t-1})}}. \tag{A14}$$

At this point, recursive forecasting can be performed by averaging over predictive results for every model using $\pi_{t+1\,|\,t,q}$ to perform the weighing across models. Therefore, DMA point predictions are given by:

$$\hat{r}^{DMA}_{t+1\,|\,t} \equiv E[r_{t+1}|r_t, r_{t-1}, \ldots, r_1] = \sum_{l=1}^{Q} \pi_{t+1\,|\,t,l}\hat{\boldsymbol{\theta}}'_t(l)\mathbf{z}_t = (\hat{\boldsymbol{\theta}}^{DMA}_{t+1\,|\,t})'\mathbf{z}_t(l), \tag{A15}$$

where $\hat{\boldsymbol{\theta}}^{DMA}_{t+1\,|\,t} \equiv \sum_{l=1}^{Q} \pi_{t+1\,|\,t,l}\hat{\boldsymbol{\theta}}_t(l)$ is the filtered, real-time forecast of the predictive regression coefficients applicable at time $t+1$ given the information available at time t.

DMS proceeds by selecting the single model with the highest value for $\pi_{t+1\,|\,t,q}$ for $q = 1, 2, \cdots, Q$ at each point in time and simply using it for forecasting:

$$\hat{r}^{DMS}_{t+1\,|\,t} = \hat{\boldsymbol{\theta}}'_t(\hat{k}_t)\mathbf{z}_t(\hat{k}_t) \quad \hat{k}_t \equiv \arg\max_{k=1,\ldots,Q} \pi_{t+1\,|\,t,k}. \tag{A16}$$

Of course, also the DMS forecast is a real-time one. In this case, we can define $\hat{\boldsymbol{\theta}}^{DMS}_{t+1\,|\,t} \equiv \hat{\boldsymbol{\theta}}_t(\hat{k}_t)$ and write about a filtered, sup-type real time forecast of the predictive regression coefficients applicable at time $t+1$ given the information available at time t.

In summary, conditional on knowing or estimating h_t for $t = 1, 2, \cdots, T$, the DMA and DMS algorithms surveyed above only involve updating formulas that identical or approximations of typical Kalman filter iterations. All recursions are simply started out by selecting a prior for $\pi_{0\,|\,0,q}$ and $\hat{\boldsymbol{\theta}}_0(q)$, $q = 1, 2, \cdots, Q$. In a very pragmatic way, we set $\pi_{0\,|\,0,q} = 1/Q$ and $\hat{\boldsymbol{\theta}}_0(q) = \mathbf{0}$ for all values of the state q, to allow the data to truly express the existence of any predictability relation.[17] As for $h_t(q)$, we follow Raftery, Kárnỳ, and Ettler (2010) and simply proceed to plug in place of $h_t(q)$ a consistent estimate that – in line with existing applied work using DMA and DMS methods – is simply based on an Exponentially Weighted Moving Average (EWMA) estimator,

$$\begin{aligned}\hat{h}_{t+1}(q) &= \sqrt{(1-\delta)\sum_{i=1}^{t+1}\delta^{i-1}(r_{t+1-i} - \hat{\boldsymbol{\theta}}'_{t-i}(q)\mathbf{z}_{t-i}(q))^2} \\ &= \sqrt{\delta\hat{h}_t(q) + (1-\delta)(r_t - \hat{\boldsymbol{\theta}}'_t(q)\mathbf{z}_t(q))^2}\end{aligned} \tag{A17}$$

with

$$\hat{h}_0(q) = T^{-1}\sum_{i=1}^{T}(r_{t+1-i} - \hat{\boldsymbol{\theta}}'_{t-i}(q)\mathbf{z}_{t-i}(q))^2.$$

Notice that, in view of our use of weekly data we set the RiskMetrics-style parameter to $\delta = 0.96$ to guarantee a not too slow decay.