The Cryptocurrency Participation Puzzle

RAN DUCHIN, DAVID H. SOLOMON, JUN TU, and XI WANG*

First Draft: August 2022
This Draft March 2023

Abstract

Ongoing zero portfolio weights in cryptocurrency are surprisingly difficult to generate in a Bayesian portfolio theory framework. With ten years of prior data, equity investors would need very pessimistic priors on mean returns to never buy cryptocurrency: -10.6% per month for Bitcoin, and -19.6% for a diversified cryptocurrency portfolio. Most priors that involve never purchasing cryptocurrency imply *shorting* it. Optimal weights are generally small, non-trivial (1-5% magnitude), frequently positive, and smooth. The certainty equivalent gains from cryptocurrency are comparable to international diversification and exceed the size anomaly. Costs (storage, fees) would need to exceed 21-39% per year to deter trading.

Keywords: Cryptocurrency, Bitcoin, Bayesian Portfolio Theory, Portfolio Choice, Non-Participation, Beliefs, Investment Frictions

JEL Codes: C11, E42, G11, G41

---

*Ran Duchin (ran.duchin@bc.edu) and David Solomon (david.solomon@bc.edu) are at Boston College Carroll School of Management, Jun Tu (tujun@smu.edu.sg) is at Singapore Management University Lee Kong Chian School of Business, and Xi Wang (wang.x@pku.edu.cn) is at Peking University School of Economics. We thank Simcha Barkai, Sam Hartzmark, Rawley Heimer, Camelia Kuhnen, and seminar participants at the University of Birmingham, Boston College, Erasmus University, the University of Georgia, the MSUFCU conference, Shanghai Jiao Tong University and Southwest Jiaotong University for helpful comments and suggestions. All remaining errors are our own.

How should one trade assets that one suspects may be a bubble? Despite its practical importance, this question is surprisingly hard to answer under existing finance theory. Part of the difficulty arises because the baseline concept of bubbles - do they even exist, and if so, how might we identify them – is itself highly contentious (e.g. Greenwood, Shleifer, and You 2019). This conceptual challenge is exacerbated by the fact that the most likely candidate assets tend to have both large run-ups in price and considerable uncertainty about their fundamentals. Such uncertainty makes it both difficult to know the chances that an asset is a bubble, and difficult to determine how investors (and thus prices) are expected to react if it is.[1]

This tension is perhaps nowhere better on display than in cryptocurrency markets. Beliefs about cryptocurrency range from it being a scam or a Ponzi scheme, to being the future of financial innovation. Importantly, these disagreements are over returns and the broad nature of the asset class, rather than fundamentals in the ordinary sense – there is no disagreement about whether Bitcoin will ever produce a US-dollar-denominated dividend. This makes it a challenge to understand pricing dynamics, since most models, including many behavioral models, rely on cash flows. No matter what an investor's beliefs are – even if he believes that the absence of cash flows for cryptocurrencies implies that their price ought to be zero – *some* portfolio decision has to be made. Investors cannot opt out of this choice simply because they dislike cryptocurrency, believe it is a bubble, or worry they don't understand it.

The answer that many investors reach is "do nothing" – 76% of US investors hold a cryptocurrency weight of zero, which we refer to as "noninvestment".[2] Such an answer represents not only a decision to not buy cryptocurrency, but also *a decision not to short it*. A zero allocation to cryptocurrency is entirely understandable as a psychological rule – "if in doubt, stick with the status quo and take no action." But is it a good economic rule? And if so, on what basis?

[1]There are many different models of bubbles. Many do not readily lend themselves to quantitative predictions as opposed to qualitative ones, especially for cryptocurrencies that lack cash flows (e.g. "If bitcoin is a bubble, when is it expected to collapse?"). See, for instance, Scheinkman and Xiong 2003, Abreu and Brunnermeier 2003 Hong and Sraer 2013, and the much larger literature on models of mispricing in general, such as Daniel, Hirshleifer, and Subrahmanyam 1998, Barberis, Shleifer, and Vishny 1998, Hong and Stein 1999 and many others.

[2]See `https://www.finder.com/how-many-people-own-cryptocurrency`

We explore this question using a Bayesian portfolio theory framework (Pástor 2000, Tu and Zhou 2010) that focuses on the role of prior beliefs in cryptocurrency market participation. Rather than taking a stand on what investors' prior beliefs should be, we offer a "conditionally normative" approach given those beliefs. We consider a relatively sophisticated investor with access to different financial markets (e.g. futures exchanges), who faces various plausible frictions. He is seeking advice on what the best tools of finance say his allocation to cryptocurrency ought to be, given his beliefs about the asset. In our framework, investors have priors about future returns. This can be thought of as a tractable reduced-form way of incorporating the predictions of different cryptocurrency models and parameters, without being restricted to a particular model of the underlying economics or psychology. For instance, models of bubbles generally imply that expected returns are negative, and these "pessimistic priors" are of particular interest.

In our base specifications, we focus on the effect of different prior beliefs about mean cryptocurrency returns, and assume that investors are approximately well-calibrated about volatility and correlations. To capture the level of certainty in these priors, we assume that investors have observed ten years of data before the series began. They update their beliefs about returns using Bayes' rule, which captures a reduced form of model uncertainty. A long enough period of high realized returns should gradually make investors more optimistic, depending on the strength of their initial beliefs. Investors initially hold the US market portfolio, and choose an optimal portfolio that adds positions in cryptocurrency, subject to reasonable levels of risk aversion.

Our main finding is that an allocation of precisely zero to cryptocurrency, year after year, is hard to generate by most pessimistic Bayesian priors. One can either believe that cryptocurrency is a bubble with large negative expected returns, or have zero portfolio allocation to cryptocurrency every period, but it is surprisingly difficult to do both. While there are good reasons to avoid a *large* allocation to cryptocurrency, it is much more difficult to justify avoiding small positive or negative portfolio weights.

We find that investors would need to be very pessimistic to never take positive weights in cryptocurrency. Prior beliefs about mean Bitcoin returns would need to be lower than -10.6% per month to justify never buying Bitcoin by the end of the sample period (February 2022). For an equally-weighted portfolio of cryptocurrency, the results are even more extreme: priors would need to be lower than -19.6% to justify never buying. These estimates, however, understate the puzzle because "never buying" is not the same as "always zero." Even if investors had priors as negative as those above, they generally imply *negative* portfolio weights. Shorting Bitcoin was possible on less reputable exchanges during the entire sample period, and short positions in futures contracts on Bitcoin began trading on the Chicago Mercantile Exchange (CME) in December 2017. We find that of priors that generate zero weights if we restrict short-selling, nearly all entail nontrivial negative weights for unconstrained investors. This is consistent with the general advice from portfolio theory that is not unique to cryptocurrency – an optimal portfolio will generally have some weight, positive or negative, in all non-redundant assets.

More broadly, we find that investors' optimal weights have four main properties - they are generally i) small, ii) nontrivial, iii) smooth, and iv) frequently positive. Across a wide range of priors about mean cryptocurrency returns (between 2% and -20% per month), optimal weights are never very large, ranging from a high of 7.3% to a low of -19.8%. Nonetheless, optimal weights are nontrivial, and most absolute weights range between 1% and 5%. The absolute value of optimal weights in Bitcoin always exceeds 0.9% at some point during the sample period for *any* ten-year prior (even including early short sales constraints), and the desired weights absent short sales constraints exceed 3.6%. As such, noninvestment is not easily explained by weights being too small to be "worth it.". Close-to-zero *average* weights over the sample period result from negative weights early in the sample and positive weights later on, rather than consistent zero weights.

Our estimates also provide evidence against the common complaint that cryptocurrency is too volatile to trade in either direction. In our setup, a *known* high volatility merely leads to lower weights, rather than zero weights. For volatility beliefs to generate weights closer to zero, investors would need to believe that cryptocurrency was even more volatile than it actually was.

While this is theoretically possible, it requires investors to have been consistently surprised over the past decade at how *smooth* cryptocurrency returns were. Similarly, cryptocurrency's high volatility doesn't lead to large changes in weights, especially over short periods – with ten years of data before the series starts, posterior means update slowly, and the high sample volatility is already factored in. As such, our conclusions are unlikely to change by adding new data.

We also compare the benefits of cryptocurrency to other assets, including those in which investors often hold zero weights (e.g., the "home bias puzzle" in international equities (French and Poterba 1991, Tesar and Werner 1995).[3] We find that under a wide range of priors, an equity market investor would perceive Bitcoin as generating certainty equivalent of returns (CER) gains that exceed those of the size (SMB) portfolio. The gains are also comparable to those from international diversification. For priors between -2% and 2% per month, perceived portfolio-level gains were 10 to 23 basis points per month, exceeding those of the MSCI world (ex-US) portfolio. At -5% priors, perceived gains are roughly half as large as international diversification. The benefits from adding cryptocurrency are thus comparable to those of other assets, for which the normative advice given by academics is that the diversification benefit is real, and investors *should* hold these assets. Failing to diversify is usually viewed as at least puzzling (e.g. French and Poterba 1991), if not a prima facie mistake, an attitude not commonly applied to cryptocurrency.

Next, we investigate whether frictions and investment costs alter our conclusions. These can include transaction costs, attention constraints, ambiguity aversion, a dislike of unfamiliar assets, storage costs, etc. We introduce a fraction-of-absolute-weight cost for investment in cryptocurrency as a reduced-form way of modeling some of the costs above. For simplicity, we consider these costs as an ongoing monthly cost, although some are likely one-off fixed costs (e.g. learning about safely storing cryptocurrency). We estimate that if costs are 20% per year or less (for either long or short weights), then there are *no* ten-year priors that result in consistently zero weights in

[3]The stock market participation literature assumes that the choice, and thus the puzzle, is between positive weights in equities and the risk-free rate. Not only does this ignore the possibility of short positions, but one could equally ask about the convertible bond participation puzzle, the mortgage-backed securities puzzle, or many other assets often allocated zero weights.

Bitcoin. For the equally weighted cryptocurrency portfolio, costs would need to exceed 30% per year. In other words, to deter investment, the implied costs of cryptocurrency investment would need to be very large, in part because the chosen weights against which costs are levied are not that large. While trading costs are often informally cited as reasons for zero weights, our results show that this claim is quantitatively difficult to support.

Which prior beliefs can justify noninvestment in cryptocurrency? Surprisingly, consistent zero weights are more likely to be achieved by a combination of being more dogmatic, but also *less* pessimistic. We find that extremely strong beliefs that cryptocurrency will earn slightly less than the risk-free rate are likely to produce consistently near-zero weights. Intuitively, such beliefs will combine with positive sample returns to generate posterior beliefs near zero, and the high volatility of returns will make it undesirable to be either long *or* short. Dogmatism will also prevent those near-zero posterior beliefs from changing over the sample period. We show that if investors are dogmatic, and presumed to have seen 50 years of past data, persistent noninvestment arises for priors of -2.4% for Bitcoin and -4.6% for the equally weighted cryptocurrency portfolio.

Our conclusions remain unchanged if investors start with different equity portfolios, such as the Fama-French factor portfolios. Beliefs that cryptocurrency is positively correlated with equity markets (rather than uncorrelated as in the baseline case) lead to moderate decreases in cutoff priors. Model uncertainty, such as robust decision-making (Hansen and Sargent 2011), leads to similar effects as beliefs in higher volatility – it shrinks desired weights, but does not significantly affect cutoff beliefs for noninvestment. We note that these analyses are not meant to exhaust all possible beliefs. For example, intuitions about bubbles might imply zero or even positive short-term expected returns, but negative expected returns at longer horizons. Implementing portfolio theory over multiple horizons is not straightforward, but if investors can construct priors over next month's returns and reevaluate, a similar myopic version of our analysis ought to apply. While our tools are assuredly not the final word on optimal portfolio allocations, they are reasonably standard, and allow us to rigorously incorporate different beliefs without relying on ad hoc justifications or informal rhetorical arguments. In the Appendix, we describe some reasons for potentially *not* having pessimistic priors. We argue that many of the puzzling aspects about Bitcoin arguably apply to gold as well, and that the parallels between the two assets deserve more consideration, even without a fully fleshed out model.

Finally, our Bayesian approach provides a different perspective on the public narrative that long positions are inherently unadvisable if one believes the asset must eventually end up at zero (whether by regulation, exchange collapse, the bubble naturally bursting, etc.). This view stems from the standard intuition that a total collapse will wipe out all of one's gains, and lead to a 100% loss on the investment. But this logic only applies to buy-and-hold positions, not to rebalanced portfolios. In our framework, target weights update slowly. Consequently, the high realized cryptocurrency returns during the sample period require rebalancing by selling cryptocurrency and buying the equity portfolio. This process effectively "locks in" many of the gains, even without this being a formal goal of the strategy. Because target weights are rarely large, a long investor whose target weight was below 5% only has a maximum downside in any single month of -5% of his portfolio. Instead, the disaster scenario for such an investor is not a 100% decline in a single month, but a 90% decline for many months in a row. In such cases, rebalancing will have the opposite effect, putting more money into a position that then continues losing. Part of this effect will be offset by the investor's posterior becoming more negative, but as before, this is a slow process given the prior certainty and the sample size. In other words, contrary to the popular narrative, for an investor who rebalances, slow collapse is much more worrisome than fast collapse.

Overall, we find the case for simultaneously believing that cryptocurrency is a bubble, and also taking weights of zero, to be much less straightforward than commonly assumed. Large positive returns ought to lead many initially pessimistic investors to eventually become optimistic. Strong beliefs that cryptocurrency is a bubble with large negative returns imply one should want to short cryptocurrency. Costs, volatility, uncertainty and other drawbacks are reasons to not take *large* weights, but much weaker reasons for avoiding even small weights. The longer cryptocurrency has a non-zero price, the stronger these arguments become. A lack of good models for long-term positive cryptocurrency prices is a license for skepticism and nonparticipation, but not an *infinite* license. It is amenable to quantification. If one resists these conclusions, our paper can be thought of as a challenge to describe what is missing in the logic we present here.

1 Related Literature

A number of existing papers study cryptocurrency returns. Y. Liu, Tsyvinski, and Wu 2022 argue that cryptocurrency market, size, and momentum capture the cross section of expected returns. Y. Liu and Tsyvinski 2021 show that user adoption of cryptocurrencies and investor attention predict cryptocurrency returns. Harvey et al. 2022 study cryptocurrency investability, realized performance, volatility, and correlation with traditional assets. Yi, Xu, and G.-J. Wang 2018, Zhang et al. 2018describe the statistical properties of cryptocurrency returns. Other papers, such as Chuen, Guo, and Y. Wang 2017, Brauneis and Mestel 2019, Flori 2019, Hrytsiuk, Babych, and Bachyshyna 2019, Rozario et al. 2020, Boiko et al. 2021, and Petukhina et al. 2021, evaluate cryptocurrency investment strategies. We complement this work by focusing explicitly on the role of beliefs. Rather than using sample returns to dispute the basis for pessimistic priors, we take such beliefs seriously and evaluate which portfolio allocations they actually justify.

Other papers attempt to model the valuation of cryptocurrency. Cong, Y. Li, and N. Wang 2021 present a dynamic model of cryptocurrency prices based on transactional demand. Biais et al. 2022 offer a general equilibrium model where the fundamental value of cryptocurrency depends on transactional benefits from future prices. Sockin and Xiong 2020 and J. Li and Mann 2018 present models of how initial coin offerings can be useful in helping to create demand for digital platforms. Pagnotta 2022 examines the security and pricing of Bitcoin in the face of potential systemic attacks. Yermack 2017 describes the general applications of blockchain technology. We add to this research by offering a Bayesian portfolio theory approach that maps investors' valuations of cryptocurrency to prior beliefs, and examines the investment behavior that they prescribe.

Yet another strand of the literature studies frictions and legal issues related to cryptocurrencies. Foley, Karlsen, and Putniņš 2019 study the illegal share of Bitcoin activity. Griffin and Shams 2020 argue that cryptocurrency prices were inflated by the supply of Tether during the 2017 boom. Makarov and Schoar 2019 study cryptocurrency price discovery and its relation to market segmentation and investor exuberance. Makarov and Schoar 2020 show that cryptocurrency markets exhibit periods of large price deviations and arbitrage opportunities. Makarov and Schoar 2021 show that Bitcoin ownership is highly concentrated and hence potentially susceptible to systemic risk. To the extent that such frictions and misuses of cryptocurrencies introduce skepticism about their value, or generate transaction costs, we complement these studies by evaluating the implications of skeptical beliefs and trading costs for cryptocurrency market participation.

Lastly, our paper adds to the literature on participation choices in different asset classes, which has mostly focused on households' non-investment in equities (e.g., Mankiw and Zeldes 1991, Campbell 2006, and Calvet, Campbell, and Sodini 2007, and many others). Unlike our paper, this literature is mostly positive in nature, trying to explain the biases or costs that might drive the observed choices. We focus on the normative aspect that is mostly just assumed for equities, namely whether investors *should* be holding the assets. Establishing this baseline is important, because it is often argued that for cryptocurrency the puzzle to be explained is the opposite one, namely why anyone *is* buying any at all. We examine optimal choices under a wide range of possible beliefs about average returns and reasonable estimates of costs and barriers to trade, and leave to future research the important question of why these are not always followed.

While our suggestion that non-participation in cryptocurrencies is *equally* puzzling as non-participation in equities is somewhat tongue in cheek, there are several serious points of comparison. First, non-participation choices are relevant for many other assets than just equities, and potential explanations for equities should be considered in terms of their application to other assets. Second, non-participation should be understood as not just failing to *buy* an asset (as in the equity literature), but failing to either buy *or short* it, if the shorting is available. Finally, cryptocurrency non-participation occurs even among investors who *are* participating in equity markets, so it is unclear ex-ante how much the same reasons should apply in both settings. We focus on the role of prior beliefs about average returns, and leave the role of other factors to future research.

2 Data

We obtain cryptocurrency prices, volumes and market capitalizations from CoinGecko, from May 1st 2013 to February 28th, 2022. This includes data on dead coins, whose returns are included in our analysis. The large volatility of prices makes it challenging to distinguish real but extreme returns from data errors. For instance, a single erroneous price can generate errors of extreme positive and negative returns on consecutive days. But this pattern can also occur in a very illiquid market if a large holder dumps a sizable volume of coins, and the price then rebounds. Similarly, days with zero volume are probably more likely to be data errors, but because turnover is correlated with returns, excluding zero volume days creates a look-ahead bias if the lack of volume is genuine. All these problems are exacerbated for small coins, whose returns are extremely volatile.

We utilize screens that exclude obvious data errors and highly illiquid coins without altering the reported return data with arbitrary changes like winsorizing. First, negative prices and market capitalizations are set to missing. Returns are dropped if:

i) Changes in market capitalization and price are incongruent: If the percentage price change and the percentage market capitalization change differ by more than 200%.[4]

ii) The market capitalization yesterday or the day before appears to be stale. That is, if the market capitalization on day t-2 and t-1 are identical, but the absolute value of the return on day t-1 exceeds 2%, returns on day t are dropped.

iii) The turnover the previous day (total volume divided by market capitalization) is less than 0.1%, or greater than 200%.

[4] For example, if the return were -50%, but the market capitalization grew by more than 150% on the same day (or vice versa), this would be dropped.

iv) Market capitalization is less than $100 million on at least three of the past five days.

The choice of lagged values for these exclusions is designed to ensure that the list of dropped coins is known before the day's returns, so there is no look-ahead bias if missing data is correlated with particular returns. We also exclude stablecoins. We generate an initial list of potential stablecoins by sorting all coins based on their maximum price deviation from $1.[5] We exclude coins whose name indicates phrases like "usd", "dollar", "stable", or other fiat currency names. For low deviation coins, and likely ambiguous cases, we use google searches to form judgments of which coins are stablecoins.

Including stablecoins has little effect on cryptocurrency returns. Winsorizing coin-level returns and market capitalizations at 0.01% or 0.1% *before* the screens are applied does not affect portfolio returns *after* adding the screens. This suggests that the screens (which largely select on stale data) eliminate the largest outliers that winsorizing would otherwise capture. Our results are similar (and slightly stronger) if the minimum market capitalization is set at $10m, but returns are much more volatile at screens of $1m.

US Market returns are taken as the CRSP value-weighted daily return from WRDS. Factor portfolios for SMB, HML and UMD are taken from Ken French's website. MSCI World Ex-US returns are taken from S&P Capital IQ. We cumulate cryptocurrency returns (which trade every day) to match equity market returns on equity-trading days. Table 1 describes the summary statistics for the main variables used in the paper.

3 Bayesian Portfolio Choice Methodology

We apply a similar framework to Pástor 2000 and Tu and Zhou 2010. There are N risky assets, and investors have priors over means, standard deviations, and correlations. The investor has access to leverage via a risk-free asset, and chooses an optimal portfolio of the risk-free asset and

[5] We do not ex-post require that the coin kept a price around $1, as a number of stablecoins have failed to maintain their peg, and just use the list to focus searches.

the estimated tangency portfolio of risky assets, given his priors and the observed data. The risky portfolio choice is generally between the equity market portfolio and cryptocurrencies, though we calculate properties for other assets used later (e.g. the size portfolio). The excess returns are $R_t = (r_{1,t}, \ldots, r_{N,t})' - r_f \mathbf{1}_{N\times1}$, which are independent and identically distributed, and follow a joint normal distribution with mean μ and covariance matrix V. To maximize expected utility, a mean-variance-utility investor chooses the optimal ω to maximize the quadratic objective function

$$U(w) = E[R_p] - \frac{\gamma}{2}Var(R_p),\ R_p = w'R_{t+1} \tag{1}$$

$$= w'\mu - \frac{\gamma}{2}w'Vw, \tag{2}$$

γ , the relative risk-aversion coefficient, is equal to three. This does not affect the estimated tangency portfolio or the Sharpe ratio, but does affect the optimal combination of the risk-free asset and the tangency portfolio, and the perceived certainty equivalent of return gains. From Markowitz 1952 , when both μ and V are known, the optimal weight is

$$w^* = \frac{1}{\gamma}V^{-1}\mu \tag{3}$$

and the maximized expected utility is

$$U(w^*) = \frac{1}{2\gamma}\mu'V^{-1}\mu = \frac{\theta^2}{2\gamma} \tag{4}$$

where $\theta^2 = \mu'V^{-1}\mu$ is the squared Sharpe ratio of the ex-ante tangency portfolio.

Calculating the optimal weights w^* in equation 3, requires the true values of μ and V, which are not known. The usual technique is to treat the sample estimates as being the true values. However, this creates a parameter uncertainty problem, and the utility under sample parameters can differ substantially from $U(w^*)$ under the true parameters.

We instead use a Bayesian method for optimal portfolio choice. The Bayesian optimal portfolio maximizes expected utility under the predictive distribution, and accounts for estimation error automatically. “Predictive distribution" refers to the fact that investors update their beliefs on the distribution of the parameters via Bayesian methods, and then use these parameter distributions to update the estimated asset return distributions. For each data realization, a different prior belief will lead to a different predictive density and a different optimal portfolio. The investor solves:

$$\max_{w} \int_{R_{t+1}} \tilde{U}(w) p\left(R_{t+1} \mid \Phi_t\right) \mathrm{d}R_{t+1} \tag{5}$$

$$= \max_{w} \int_{R_{t+1}} \int_{\mu} \int_{V} \tilde{U}(w) p\left(R_{t+1}, \mu, V \mid \Phi_t\right) \mathrm{d}\mu \mathrm{d}V \; \mathrm{d}R_{t+1}, \tag{6}$$

where $\tilde{U}(w)$ stands for the utility of holding a portfolio w at time t+1, $p\left(R_{t+1} \mid \Phi_t\right)$ is the predictive density, Φ_t denotes the data available at time t, and

$$p\left(R_{t+1}, \mu, V \mid \Phi_t\right) = p\left(R_{t+1} \mid \mu, V, \Phi_t\right) p\left(\mu, V \mid \Phi_t\right) \tag{7}$$

where $p\left(\mu, V \mid \Phi_T\right)$ is the posterior density of μ and V.

Let $\mu*$ and V^* stand for the first two moments of the predictive density $p\left(R_{t+1} \mid \Phi_t\right)$. The optimal weight is given by,

$$w_{Bayes} = \frac{1}{\gamma}(V^*)^{-1}\mu^* \tag{8}$$

Equation (8) shows intuitively that the weight in a given asset increases with its predicted mean. It also increases as the asset’s contribution to portfolio variance decreases (i.e. its own variance decreases, or its correlation with other assets decreases). Weights on risky assets also decrease as the risk aversion coefficient increases, but this only occurs through a change in leverage.

The expected utility or ex-ante certainty-equivalent return of using w_{Bayes} is given by

$$EU_{Bayes} = w'_{Bayes}\mu^* - \frac{\gamma}{2} w'_{Bayes} V^* w_{Bayes} \tag{9}$$

We solve the restricted optimization problem

$$\max_{w_{res}} \int_{R_{t+1}} \int_{\mu} \int_{V} \tilde{U}(w) p\left(R_{t+1}, \mu, V \mid \Phi_t\right) \mathrm{d}\mu \mathrm{d}V \, \mathrm{d}R_{t+1},$$

to obtain the allocation w_{res}, the optimal weight conditional on no cryptocurrency investment ($w_{crypto} = 0$). The expected utility or ex-ante certainty-equivalent return of w_{res} is,

$$EU_{res} = w_{res}' \mu^* - \frac{\gamma}{2} w_{res}' V^* w_{res} \tag{10}$$

This expected utility is evaluated based on the same μ^* and V^* from the predictive density.

We can define the difference between the unrestricted Bayesian expected utility (EU_{Bayes}) and the expected utility with cryptocurrency weights restricted to zero (EU_{res}) as the ex-ante or perceived certainty-equivalent return (CER) gain from cryptocurrency.

Similarly, we measure the ex-post gain from investing in cryptocurrency for each prior. To do this, we: (1) calculate the portfolio weights of w_{Bayes} and w_{res} period by period; (2) compute the ex-post returns of these two portfolios ; (3) compute the mean $\hat{\mu}_{bayes}$, $\hat{\mu}_{res}$ and variance $\hat{\sigma}^2_{bayes}$, $\hat{\sigma}^2_{res}$ of the realized portfolio return sequences. Then, the ex-post certainty-equivalent return (CER) is

$$U_{Bayes} = \hat{\mu}_{bayes} - \frac{\gamma}{2} \hat{\sigma}^2_{bayes}$$
$$U_{res} = \hat{\mu}_{res} - \frac{\gamma}{2} \hat{\sigma}^2_{res}$$

The difference between these two terms (U_{Bayes} - U_{res}) is the ex-post certainty-equivalent return (CER) gain from investing in cryptocurrency. As weights w here represents the investment weight in risky assets, we can define leverage as $\sum_i^N w_i$. If it is larger than 1, an investor has to borrow to invest in risky assets. If less than 1, an investor invests both in risky assets and the risk-free asset.

Without imposing diffuse priors on μ and V, we have no analytical form for μ^* and V^*. We thus estimate them by an MCMC algorithm. As μ^* and V^* vary over time, we estimate them each period using a Gibbs sampling method, described in the Internet Appendix.[6]

4 Results

4.1 Prior Beliefs and Cryptocurrency Participation

We first compute how pessimistic an investor's prior beliefs about cryptocurrency would need to be to justify non-participation over time. For each posterior belief about covariance, there is a single posterior belief about mean excess returns that corresponds to precisely zero investment. All higher posterior beliefs result in a desire to buy the asset, and all lower beliefs create a desire to short the asset. Unless otherwise noted, statements about returns reference excess returns. Solving for these cutoff, zero-weight prior beliefs about mean excess returns requires several parameters: (1) The initial assets that the investor holds; (2) A choice of cryptocurrency; (3) A strength or dogmatism of prior beliefs; and (4) A time period when the evaluation is made. We begin with a baseline specification of these parameters, which we vary later in the paper.

We assume that the investor starts with the CRSP value-weighted market portfolio as his risky asset. This captures someone following the frequently-dispensed advice to hold diversified equity index funds. We consider three cryptocurrency portfolios: Bitcoin-only, as well as equal-weighted and value-weighted portfolios of coins (with a prior market capitalization greater than $100m). For the strength of beliefs, we assume that an investor updates as if he observed ten years of monthly data with the parameters in question, and adds each new month's returns to work out the posterior values. We assume beliefs about variance that roughly correspond to ex-post sample outcomes, with Bitcoin having a variance 150 times the market (ex-post was 143), the value-weighted cryptocurrency portfolio having a variance 170 times the market (ex-post was 166), and the equal-weighted portfolio having a variance 625 times the market (ex-post was 612). Lastly, we assume that the investor believes cryptocurrency to be uncorrelated with the market.

[6] Available at `https://tinyurl.com/2p8njmjj`

In Table 2, we examine the cutoff prior mean beliefs that correspond to zero investment. Panel A presents the prior mean beliefs that, at each point in time, map to zero weights. The first row focuses on Bitcoin. Even early in the sample, quite large pessimistic priors about mean excess returns were needed to justify non-investment: -4.7% per month in December 2013, with similar numbers in 2014 and 2015. However, the large returns in 2017 significantly changed this, with zero investment priors dropping to -8.5% per month in December 2017, and finishing the sample period at -10.3% in February 2022.

Row 2 repeats the analysis for the equal-weighted cryptocurrency portfolio. Because its mean returns are higher than Bitcoin, prior beliefs would have needed to be even more pessimistic, from -9.1% per month in 2014, to -19.2% per month at the end of the sample. Row 3 studies the value-weighted cryptocurrency portfolio. Cutoff priors lie between the Bitcoin-only and the equal-weighted portfolio, due to the large weight of Bitcoin in the value-weighted portfolio. This pattern is consistent throughout our analyses.

In Panel B, we consider what prior beliefs would lead investors to *never* take a positive weight up to each point in time. Unlike in Panel A, the required prior mean strictly ratchets downwards over time. The priors needed to never buy cryptocurrency are similar but slightly lower than in Panel A. The final end-of-sample values are -10.6% for Bitcoin, -19.6% for the equal-weighted cryptocurrency portfolio, and -11.0% for the value-weighted portfolio. In Figure 1, we plot the continuous time series of prior beliefs required for non-investment. In the Internet Appendix, we show that the results at local price minima and maxima are similar to the annual snapshots.

It is worth emphasizing just how large these numbers are. On a compounded basis, the annual expected losses would need to be roughly -70% per year for Bitcoin, and -90% per year for the equal-weighted portfolio to justify never buying throughout the sample period. Even pessimistic investors who expected cryptocurrency to lose, say, 50% per year at the beginning of the sample, would nonetheless have taken a long position at some point when the posterior means become positive due to the high realized returns.

4.2 Short Sales Constraints and Optimal Cryptocurrency Weights

Short sales constraints are a large potential driver of zero portfolio weights, because when investors cannot short sell, beliefs implying negative desired weights lead to zero realized weights. For Bitcoin, the ability to short sell has a clear break on December 10th, 2017, when the CBOE launched Bitcoin futures (followed by the CME shortly after on December 17th, 2017). This made negative weights in Bitcoin possible even for US investors with only access to traditional, regulated financial markets. Shorting on less reputable exchanges was possible throughout our sample period.[7] While shorting is less straightforward for other cryptocurrencies, the ability to short has generally increased over time, making non-participation later in the sample more puzzling.

While the results in Table 2 show the difficulty of justifying zero weights year after year, it is important to understand *how much* one should invest in cryptocurrency. If optimal weights are very close to zero, these may disappear when plausible frictions are introduced, or if investors have a heuristic that small trades are "not worth it". Since cryptocurrency has had large and volatile returns, small weights may still have considerable effects, so the heuristic is not obviously correct, despite its plausibility. We calculate cryptocurrency portfolio weights for a wide range of opinions about cryptocurrency, trying to capture the range of commonly expressed opinions. These include mildly optimistic priors (2% and 1% per month), pessimistic priors (-1%, -2%, -5%, -10% and -20%), neutral (no risk-exposure, market-efficiency) priors of 0% per month, and flat/diffuse priors where the investor believes in the sample mean at each point. While one could study highly optimistic investors, we suspect (anecdotally) that most optimists actually held cryptocurrency, and so do not have a participation puzzle.

Table 3 summarizes the range of optimal cryptocurrency portfolio weights, where each row corresponds to a different prior belief. For brevity, we only report the results for Bitcoin and the equal-weighted cryptocurrency portfolio in the rest of the paper, and include the main results for the value-weighted portfolio in the Internet Appendix.

[7] See, for instance, this Quartz article from 2013 describing how to short Bitcoin on Bitfinex and ICBIT at that time: `https://qz.com/69630/how-to-short-Bitcoins-if-you-really-must/`.

Panel A of Table 3 summarizes the range of optimal Bitcoin portfolio weights. Average weights range from 5.9% for the 2% prior to -9.6% for the -20% prior, in Column 1. The largest absolute weights result from diffuse priors, where the average weight is 16.0% (though, as we will see later, these do not lead to ex-post desirable portfolios). Columns 2 and 3 show maximum and minimum weights over the sample. These are obviously more extreme than average weights, but the difference is not that large, as extreme weights range from a maximum of 7.3% for 2% priors, and a minimum of -19.8% for -20% priors. Column 4 shows final (end-of-sample) weights, which are close to their maximum values. All priors between -10% and 2% show positive final weights (ranging from 0.2% to 7.0%).

Column 5 shows the fraction of months with positive weights. Priors of 0 and above are essentially always positive, but even priors of -1% and -2% have positive weights in close to 95% of the months. Even quite pessimistic priors such as -5% end up having positive weights over half the time (though extremely pessimistic priors of -10% and -20% result in always or nearly always being short). In other words, short sales constraints are unlikely to explain non-participation for most priors considered.

Columns 6 through 9 show the fraction of months with absolute weights above 0.5%, 1%, 2% and 5% respectively. If absolute weights above 2% are considered "worth it", this covers over 95% of months for priors of -1 or above, and 79% of the months for priors of -2%. -5% and -10% priors have both slightly less than half of months above 2% absolute weights, though at -20% priors all months exceed 2%. At lower thresholds for meaningful investment, such as 0.5% or 1% weights, most months for most priors meet the threshold. Weights above 5% are generally only reached for optimistic priors or extremely pessimistic priors. Column 10 shows the first date weights are positive, while columns and 11 and 12 show the mean and standard deviation of leverage choices. Optimists are somewhat more levered than pessimists, but the differences are not large.

Importantly, desired weights in Bitcoin are not that large in either direction (mostly in the 1-5% range), even under a wide range of priors. This reflects the fact that cryptocurrency is very volatile, and investors are assumed to be well calibrated on this aspect. This assumption also seems reasonable, as both cryptocurrency boosters and skeptics tend to agree on its high volatility, and second moments are generally easier to estimate than first moments. Being correctly calibrated on cryptocurrency volatility (as investors are here) goes a long way towards preventing investors "blowing themselves up", even under very different prior means. We return to this issue later when we consider the ex-post effects of cryptocurrency. The high volatility of cryptocurrency is a strong reason to only take *small* positions, but this is not the same as taking zero positions.

Panel B studies the equal-weighted cryptocurrency portfolio, and shows similar patterns. The higher average returns of the equal-weighted cryptocurrency portfolio compared to Bitcoin leads to more priors with a sizable fraction of positive weights. Even fairly pessimistic priors of -5% for the equal-weighted portfolio resulted in 94% of months with desired long positions. By contrast, the higher volatility of the equal-weighted portfolio means that the magnitude of weights chosen is generally smaller than for Bitcoin.

Figure 2 shows how weights evolve over time for each prior, for Bitcoin (Panel A) and equal-weighted cryptocurrency (Panel B). Despite the volatility of cryptocurrency, optimal weights changed relatively smoothly for all the priors. The only exception, not graphed, is diffuse priors, whose weights are much more volatile. This reinforces the generally useful role of Bayesian methods for portfolio theory - informative priors are able to prevent large swings in beliefs with each new data point, especially at the beginning of the series when the available sample data is limited. This also implies that the addition of new cryptocurrency returns is unlikely to alter our conclusions in the medium term. It would require a prolonged period of very negative returns, enough to offset both the informative priors imposed and the existing sample data already incorporated. Some results, like the "priors required to never buy", are a ratchet that can only get more negative over time, regardless of future data. While cryptocurrency monthly returns are volatile, posteriors about means change relatively slowly once an informative prior is imposed, even more so when additional data has been observed. As long as the volatility of returns is reasonably stable (and understood), optimal actions do not change rapidly.

Panels C and D of Figure 2 show how maximum desired weights vary with priors. For each prior we compute the maximum absolute weight that an investor desires over the sample (absent short sale constraints). In Panel C, all investors will want an absolute weight in Bitcoin of at least 3.7% at some point in the sample period (with the minimum corresponding to a prior of -0.041, and all other priors desiring larger weights at some point). In Panel D, for equal-weighted cryptocurrency, all investors would want a weight of at least 1.6% at some point (with the minimum corresponding to a prior of -0.077). Both graphs have a v-shape, so desired maximum weights increase as investors become more pessimistic or optimistic.

Panels E and F repeat this analysis with short sales constraints. Panel E shows that for Bitcoin (with no shorting before 2017), all ten-year-prior investors desire an absolute weight of at least 0.9% at some point (with the minimum corresponding to a prior of -0.089). Panel F shows that for the equal-weighted cryptocurrency portfolio (which we assume cannot be shorted throughout the sample period), investors with all priors above -0.196 desire a positive weight at some point.

Overall, a wide range of priors map to positive cryptocurrency weights at some point in the sample. Non-participation by such investors is not easily explained by short sales constraints. Very pessimistic priors map to consistent negative weights in cryptocurrency, but such investors should have shorted Bitcoin once this became easier later in the sample. While the weights in question are small, befitting a volatile investment, they are also nontrivial, often in the 1-5% range. Despite this volatility, weights are surprisingly stable over the sample, implying that new data is unlikely to rapidly alter these results.

4.3 Quantifying the Perceived Benefits of Cryptocurrency Participation

An alternative way to measure the desirability of cryptocurrency investments is to estimate the expected benefits that they provide. We do this by calculating the certainty equivalent benefit of adding cryptocurrencies to investors' existing portfolios. We assume a constant relative risk aversion of 3, which allows us to convert the distribution of risky returns to a certainty equivalent return (CER) - the investor's constant return equivalent to the risky return. We calculate the marginal value of cryptocurrency as the difference between the CER of the baseline market portfolio excluding cryptocurrencies and the CER of the optimal portfolio that combines the market portfolio and cryptocurrencies.

We present these results in Table 4. Panel A considers Bitcoin whereas Panel B considers the equal-weighted cryptocurrency portfolio. The columns correspond to the year in question, while the rows correspond to the same range of priors considered before. Finally, the colors represent the direction of the position. Black values correspond to CERs generated by short positions, whereas green values correspond to long positions.

We find that cryptocurrency investments produce sizable certainty equivalents of returns over time and across priors. The final column, for the end of the sample period, shows that among positive weight positions (in green) the perceived gains decrease as priors become more negative. Perceived CER gains per month are 23 b.p. for 2% priors, 19 b.p. for 1% priors, 16 b.p. for 0% priors, 13 b.p. for -1% priors, 10 b.p. for -2% priors, and 4 b.p. for -5% priors. Once priors cross the cutoff that produces zero end-of-sample desired weights (-10.3%, with a CER of zero), the estimated benefits start rising again. Thus, at priors of -10%, investors perceive (very small) benefits to buying Bitcoin at roughly 0 b.p., and at priors of -20% they perceive gains of 14 b.p. from shorting. Consistent with the extreme desired weights in Table 3, flat priors provide large perceived gains throughout the sample period. Lastly, consistent with our previous findings, the fraction of CER gains generated by long positions increases over time.

Importantly, the above CER estimates do not imply that investors were actually right, nor that they did (or will) earn such gains. For instance, -20% priors produce very high ex-ante CERs throughout the sample despite the disastrous performance that shorting Bitcoin since 2013 would have produced. Instead, they show that, given a set of beliefs, investors *perceived* this gain from longing or shorting cryptocurrencies.

Panel B presents the same CER estimates for the equal-weighted cryptocurrency portfolio. The broad patterns are similar to those of Bitcoin only, with CERs ranging from 16 b.p. for priors of 2% down to 3 b.p. for priors of -10%. Figure 3 graphs end-of-sample CERs versus priors. The graphs show a U-shape, but it is fairly flat over most likely priors. This is consistent with the snapshots in the Table, which show a surprising consistency in perceived benefits at the end of the sample period for investors who started off with very different priors. In the Internet Appendix, we show the gains in terms of Sharpe Ratios, which follow a similar pattern.

4.4 Comparison of Cryptocurrency Benefits with Other Assets

The CER estimates in Table 4 can also be interpreted as the portfolio-level amount that investors would be willing to pay per month to access the assets. For instance, investors with a 2% prior would be willing to pay 23 basis points per month to access Bitcoin. This is comparable to the fees on a very expensive mutual fund, but levied on the *entire portfolio*, not just the small cryptocurrency portion. Next, we provide another comparison - between the perceived benefits (CERs) of cryptocurrencies and those of other assets. We aim to explore whether the CER estimates are specific to cryptocurrencies or merely reflect the benefits of diversifying to assets other than the value-weighted market portfolio.

We repeat the analyses in Table 4 for several prominent equity portfolios. We start with the CER of the market portfolio alone, and consider the end-of-sample increase in CER from adding the new assets. We examine the MSCI world ex-US portfolio, and the size (SMB), value (HML) and momentum (UMD) anomaly portfolios (Fama and French 1993, Carhart 1997). These not only represent important variables for explaining the cross-section of returns (Fama and French 1993), but are also important asset classes in the ETF space, and for value and size, the basis of Morningstar fund classifications. For these portfolios, we imagine investors who approximately follow academic finance consensus wisdom. As such, their beliefs about portfolio returns come from the time-series of observed returns, without imposing ex-ante priors. For instance, we assume that investors began with diffuse priors about the SMB portfolio at the beginning of the sample period (1926), and observed its returns each month from 1926 until the end of the sample period.

In Table 5 we compare the end-of-sample CER gains from these investments to the CER gains from cryptocurrencies . Because we are interested in the relative benefit of each asset, for ease of comparison we scale each CER by the CER obtained from adding the MSCI world ex-US portfolio (10.4 b.p.). This captures the benefits of international diversification, and represents an intuitive benchmark, where unlike cryptocurrencies, traditional advice is that one should hold these assets.

The ex-ante gains from investing in Bitcoin exceed those from the MSCI World Ex-US (i.e. the ratio is greater than one) for priors above -1%, ranging from 1.24 times at a prior of -1%, to 2.2 times at a prior of 2%. At -5% priors, the gains are roughly a third of international diversification, and while they are very small for -10% priors, by -20% priors they are large again. Access to Bitcoin exceeded the perceived benefits of SMB over a wide range of priors (i.e. the "Bitcoin to MSCI" ratio exceeds the "SMB to MSCI" ratio), though not HML and UMD . For equal-weighted cryptocurrency, the ratios are similar. These results show that the perceived gains of cryptocurrency are comparable to several portfolios that the academic literature has considered important.

4.5 Investment Costs

Next, we examine how costs may deter investors from taking non-zero positions. One such cost is ambiguity aversion (Epstein and T. Wang 1994). Investors may have ambiguity or entry costs from understanding cryptocurrency technology and its value proposition (e.g., Blockchain technology, ledger storage, public and private key cryptography). Here, we consider such costs as a general disutility from investing in unfamiliar assets, and later examine more formal treatments of model uncertainty. Another cost is the safe storage of cryptocurrencies. Most financial assets have extensive avenues to recover assets where credentials have been lost or stolen. Cryptocurrencies present many challenges in this regard - assets on the Blockchain are gone irretrievably if private keys are lost (and likely also if they are stolen). Passing assets onto heirs, while also ensuring that said heirs cannot steal them while the owner is alive, is also technically challenging. Finally, there are more basic transaction costs, including bid-ask spreads, impact, and exchange fees.

While some of these costs are fixed (e.g. learning about the mechanics of investment), others are likely ongoing (e.g. losses from storage). For simplicity, we model costs as ongoing costs (Vissing-Jorgensen 2003, Fagereng, Gottlieb, and Guiso 2017, Briggs et al. 2021), and leave the study of one-off costs for future research.[8] (F. Gomes and Michaelides 2005, Haliassos and Michaelides 2003, Abel, Eberly, and Panageas 2013) These costs are a simple way to model various costs of trading cryptocurrencies, which may make investors choose a zero weight even when the costless version of portfolio theory implies a non-zero weight. For simplicity, we model these costs as being symmetric for both positive and negative weights, but many of the costs (e.g. risk of theft) are larger for long positions than for short ones. The magnitude of some of these costs is not straightforward - psychological costs like ambiguity aversion, for instance, but also spreads and market impact early in Bitcoin's trading history (or for obscure coins today). Instead we solve for how large the costs would have to be to deter investment, given a set of beliefs. If the CER gain from adding cryptocurrency is greater than the costs, investors should add it to their portfolios.

In Table 6, we provide two sets of analyses to address this question. In Panels A and B, we fix a set of ongoing costs, and solve for the range of prior means that would map to zero weights throughout the sample period. In Panels C and D, we reverse these calculations, and solve for the costs that would make an investor choose a weight of zero given a prior mean belief. These costs are applied as a fraction of the absolute value of the position in cryptocurrency (e.g., the ambiguity aversion cost is proportional to how many dollars one invests long or short in cryptocurrency). We consider annual costs of 10%, 15%, 20%, 30%, and 50% of absolute weights. As percentages, these are enormously higher than most equivalent assets, and thus represent conservative estimates.

Panels A (for Bitcoin) and B (for equal-weighted cryptocurrency) take a measure of costs, and show the range of prior beliefs that correspond to non-investment in cryptocurrencies up to the

[8] In terms of magnitudes, one can consider the present value of ongoing costs as approximating the fixed cost form, though the portfolio implications are not identical.

end of each sample year. We find that there are *no* priors that map to consistent non-investment under the relatively "low" costs of 10% per year. By the end of the sample period, there are also no priors that map to non-investment at 15% or 20% costs. For the equal-weighted cryptocurrency portfolio, the estimates are more extreme: 20% costs are insufficient to deter trade for any priors, and even 30% costs do not deter trading for any priors by the end of the sample period.

Panels C and D reverse the calculation - for each prior, we solve for the minimum cost that would justify non-investment for all annual snapshots up to each sample year. This panel, unlike A and B, applies our regular assumptions about short sales constraints. For Bitcoin, once shorting is allowed in 2017, the emerging pattern is U-shaped: Investors with -10% priors have the lowest required costs (9.3% annual costs, at the end of the sample), compared with costs of 42.7% for 2% priors and 48.2% for -20% priors. These findings are consistent with the earlier results that extreme priors map to extreme desired weights, and therefore require higher costs to deter investment. As in Panels A and B, these costs tend to increase towards the end of the sample period.

It is worth emphasizing just how large these estimates are. Costs that deter investment *start* at over 20% per year for Bitcoin and 39% for the equal-weighted cryptocurrency portfolio. This mitigates concerns about the exact form of the costs, as the costs required to deter investment are large to the point of implausibility. Lastly, in the Internet Appendix we also study the effect of annual investment costs on certainty equivalent return (CER) gains from cryptocurrency.

4.6 Ex-Post Benefits of Investing in Cryptocurrency

Up to now, we have calculated the certainty equivalents of returns (CER) on an ex-ante basis. That is, at each point we computed how investors would have perceived cryptocurrency at the time, given their prior beliefs. This is separate from whether the trades actually performed well. Investors with pessimistic priors would have initially believed that shorting Bitcoin would be very profitable, and the most pessimistic retained that view at the end of the sample. Given Bitcoin's high returns during the sample, ex-post they were likely to have been disappointed.

We assess the ex-post performance on a distributional basis. Investors assume that the distribution of portfolio returns they have received up to that point (given their time-varying choice of weights) were to continue indefinitely, and assess whether they would have preferred this distribution to that of the equity market portfolio alone. This differs from a pure test of whether cryptocurrency beat equities, as it also takes into account volatility. If an investor's beliefs are too optimistic before a period where Bitcoin has high returns, he can nonetheless end up preferring the market on a distributional basis because betting too heavily on Bitcoin exposed him ex-post to higher volatility than he would like.

In Table 7, we calculate investors' CER for the ex-post distribution of portfolio returns relative to the ex-post distribution of the equity market portfolio alone. This is done at annual snapshots for various priors. Finally, we also compute the maximum possible ex-post gain and the prior beliefs that correspond to this maximum.

Panel A examines Bitcoin. At the end of the sample, Bitcoin had positive ex-post CERs for all priors above -2%, from 0.41% at 2% priors to 0.13% for priors of -2%. In other words, investors who thought that Bitcoin would lose 2% per month, or almost 22% per year, nonetheless ended up happy with their portfolio, as they switched to positive weights in November 2014 (from Table 3). The most pessimistic priors were, unsurprisingly, very unhappy with their ex-post performance, reaching a CER of -2.35% per month for -20% priors. Importantly, gains do not increase monotonically with optimism. The maximum ex-post CER gain was 0.65% per month, from a prior of 11.2%. While more optimistic investors would have taken larger weights and made higher returns, they also would have experienced more volatility, and their overall CER gain was actually lower.

Relative to the ex-ante CERs in Table 4, there is more time-series variation in how the portfolios were perceived ex-post for a given prior. The big shift occurred in 2017, where both optimistic and mildly pessimistic priors began to show large ex-post gains. Panel B shows similar patterns for the equal-weighted cryptocurrency portfolio.

Figure 4 plots how end-of-sample ex-post CERs vary with initial priors. The figure confirms the intuition from the table above - mild pessimists ended up happy ex-post (because they changed their posterior beliefs to optimism as the sample progressed), and mild to moderate optimists ended up happier still. While it was possible to be *so* optimistic as to actually have negative CERs, these only occur at enormous priors of around 25% per month. It is tempting to assume that the lesson here is that it was hard to be too optimistic about cryptocurrency ex-post. However, it is important to remember the importance of being well-calibrated about volatility, which prevents excessive optimism translating into enormous portfolio weights. It also highlights the strengths of the Bayesian framework in calculating reasonable weights under very different mean beliefs.

In Figure 5, we combine the ex-ante and ex-post assessments into a single graph. We can conceive of three dimensions of investor behavior at each point in time:

1. Were they on average long or short beforehand?
2. Are they long or short at that point?
3. Are they happy or unhappy ex-post with the returns they have received?

We use a combination of color and shading to represent these dimensions graphically. Green denotes being long both beforehand and currently, red denotes being short both beforehand and currently, and yellow denotes being short beforehand but currently long. The final possibility (long beforehand but now short) does not occur in the data. Shaded regions are happy ex-post at that point, and unshaded regions are unhappy ex-post.

We find that most pessimistic regions ended up ex-post unhappy (i.e. unshaded). Among initial pessimists, only the mildly pessimistic ended up ex-post happy, due to switching to positive weights early in the sample period. All investors who were on average short at the end of the sample period were ex-post unhappy, even those that had switched to long holdings by the end (the yellow bars). Some of the "on average long" priors (green bars) still were ex-post unhappy (unshaded), primarily those who switched to long positions relatively late. Extremely pessimistic investors stayed short the whole time, but ended up even more unhappy ex-post. Finally, the most extremely optimistic priors were ex-post unhappy, but only for enormously optimistic beliefs. Panel B shows that for the equal-weighted cryptocurrency portfolio, a wider range of priors resulted in investors being long (on average) and ex-post happy by the end of the sample period.

4.7 Robustness and Extensions

Finally, in Table 8 we vary our main specification, changing the strength of prior beliefs, baseline asset portfolios, beliefs about correlations and volatility, dropping some years' data, and adding model uncertainty. First, we examine how the strength of one's priors (i.e. the level of certainty or dogmatism, not the prior mean) changes the results. It is unclear what level of dogmatism about cryptocurrency returns is reasonable. We explore values from having seen three or five years worth of data, to 30 or 50 years of data (with ten years being the baseline in Table 2). Panel A examines the cutoff priors that map to zero investment under different dogmatism.

As expected, greater dogmatism leads to less negative cutoff prior means. At the end of the sample, a prior with three years of data would require a cutoff mean of -34.1% for Bitcoin, while dogmatism equivalent to 30 or 50 years of data maps to cutoff priors of -3.8% and -2.4% respectively. For the equal-weighted cryptocurrency portfolio, three year priors require a belief of -62.9% per month, while 30 and 50 year priors require means of -7.2% and -4.6% per month, respectively. While not tabulated, more dogmatic priors also predictably lead to more stable cutoff weights over different years of the sample.

Overall, more dogmatic beliefs are more likely to lead to consistent zero weights than more pessimistic beliefs. As priors become sufficiently strong and posterior beliefs shrink to zero, desired portfolio weights and perceived certainty equivalent gains also shrink. In the limit, increasing dogmatism can justify almost any behavior, as it places less and less weight on the observed data. We note, however, that this version of extreme dogmatism does not easily explain the rhetoric of committed cryptocurrency skeptics with zero weights. At higher dogmatism levels, the means required for non-investment are those of a slowly deflating bubble, earning slightly less than the risk-free rate. At extreme dogmatism levels, required beliefs converge on priors of market efficiency under zero risk exposure, where cryptocurrency earns the risk-free rate. These beliefs do not easily map to rhetoric that cryptocurrency is a bubble about to burst.

In Panel B, we examine the effect of different initial assets in addition to the value-weighted market portfolio. The priors required for non-investment are very similar regardless of whether one also holds SMB and HML (row 2), SMB, HML and UMD (row 3), or SMB, HML, UMD and MSCI World ex-US. These results occur because the additional assets are broadly uncorrelated with cryptocurrency returns, both under investors' priors (assumed as zero), and in the data (Table 1. This suggests that adding further equity portfolios is unlikely to change the results.

In Panel C, we vary investors' priors over the correlation of cryptocurrency with equity markets, and solve for the end-of-sample zero investment cutoff mean. Beliefs in higher correlations have nontrivial effects on cutoff means. For Bitcoin, correlations of 0.1, 0.2 and 0.3 produce cutoff means of -9.1%, -7.8% and -6.5% respectively (relative to a zero correlation baseline of -10.6%). For the equal-weighted cryptocurrency portfolio, cutoff means are -16.3%, -13.7%, and -11.1% respectively (relative to a baseline of -19.6%).

Panels D and E explore the effects of different prior beliefs about volatility. Volatility does not affect zero-investment cutoff beliefs, but only weights on either side. We examine how prior beliefs about volatility affect desired weights, for a range of beliefs about means. We consider priors about volatility ranging from 0.2 and 0.5 times sample values, to 2 and 5 times sample values. We show the results for Bitcoin, with the results for the equal-weighted and value-weighted cryptocurrency portfolios in the Internet Appendix. Panel D presents average weights, and Panel E presents end-of-sample weights. Both panels show a large effect whereby beliefs in higher volatility shrink weights towards zero. If volatility beliefs are twice the sample average, then average weights are roughly between half and a quarter as large, depending on priors. At 5 times sample averages, they are generally less than a tenth as large. The effects on end-of-sample weights are similar. Overall, beliefs in higher volatility do not on their own easily produce zero weights in a standard framework, but they can amplify the effect of other frictions by making desired weights much smaller. In the Internet Appendix, we also provide estimates of the CER gains from investing in cryptocurrencies for different volatility priors.

Panel F examines cutoff beliefs if early data is ignored as being "unrepresentative". Early cryptocurrency returns may be downweighted if they were to be viewed as unsustainable in the long run. We consider the simplest version, whereby priors are only combined with data later than a certain point, and solve for the cutoff end-of-sample beliefs required for never investing. We find that removing 2013 data somewhat lowers required beliefs, and removing data up to 2017 lowers them further. On the other hand, removing more recent periods with high realized returns has the opposite effects. These methods are extremely ad hoc, as they begin with a sophisticated Bayesian updating process and arbitrarily downweight certain returns to zero. A full version of updating under priors of non-I.I.D. returns is a worthwhile challenge, but beyond the scope of this paper.

Finally, we consider the objection that people don't invest in cryptocurrency because "they don't understand it". This statement seems intuitive, but is not straightforward to interpret. In the simplest Bayesian framework, "not understanding something" maps to diffuse priors. But since this would lead to people just weighting sample returns, and hence enormous positions, this seems to not be what is meant. In Panels G and H we examine one alternative way to formalize a lack of understanding, namely model uncertainty. Following Hansen and Sargent 2001, Anderson, Hansen, and Sargent 2003 and Anderson and Cheng 2016, investors have a form of "worst case scenario" beliefs about model uncertainty, and solve a robust portfolio choice problem. After they make their choice, an adversarial agent pays a cost to perturb the probability distribution of returns in a maximally costly way for investors' utility. Investors' concerns about model misspecification, measured by ambiguity aversion, are equivalent to the adversarial agent's perturbation cost. Investors choose portfolios that maximize their utility, considering the adversarial agent's attempt to minimize their utility.[9] Details are in the Internet Appendix. Model uncertainty ends up operating similar to belief in higher volatility. Panels G and H present the results for Bitcoin, with

[9]Model uncertainty and robust portfolio choice can also be considered a model of superstition about being maximally unlucky, rather than a rational framework.

the results for the equal-weighted and value-weighted cryptocurrency portfolios in the Internet Appendix. In Panel G, cutoff beliefs are affected very little by an ambiguity aversion coefficient of 4 (Anderson and Cheng 2016). Instead, desired weights are roughly a third to a half as large (Panel H). While model uncertainty can reduce weights, it does not have any special role for zero.

There are other ways to model a lack of understanding. Investors may feel “certain that they are uncertain”, having implicitly seen lots of data from (and thus have very strong priors about) a very high standard deviation distribution (so sample data does not change their beliefs much). Bayesian frameworks do not sharply distinguish between parameter uncertainty and model uncertainty, and such beliefs resemble strong beliefs in high volatility. Alternatively, dislike of uncertainty may be a cost that reduces utility directly. This will struggle for the same reason as costs explanations, unless the disutility is fixed and large for the first dollar of investment in either direction. As noted earlier, one can arrive at zero by strong default of “if in doubt, do nothing”, with zero being the default. This is intuitively plausible, but it seems much closer to a behavioral heuristic than a fully-fleshed-out rational expectations model. Understanding exactly how “uncertainty over cryptocurrency” should be thought of is an interesting question for future research.

5 Alternative Modeling Choices

Our analysis illustrates reasonable ways to think about cryptocurrency returns using Bayesian portfolio theory, but is not meant to be exhaustive. There are other alternatives that require greater modifications than we have used so far. One is models that predict different conditional expected returns over time. One could imagine conditional beliefs such as "the more a bubble inflates, the higher the chance of crashing", or "once cryptocurrency becomes large enough, it cannot grow at the same rate as before". Portfolio theory generally is both static and single-period. It is an interesting question, but beyond the scope of the paper, as to how to model belief updating in a dynamic setting if returns are not believed to be independent and identically distributed. Nonetheless, it seems possible to reduce complex multi-period beliefs to single-period versions that are updated each period. As long as one can form a posterior distribution over next month's returns, one can solve the one-period problem, though next month one's beliefs will change by some different updating rule. Such versions will be myopic, forecasting only one period at a time, rather than optimizing for the full path of future expectations. The iterated one-period version may be a decent approximation of some versions of multi-period beliefs, though it falls short of a full model.

Another possibility is that investors have some two-part belief process, combining a binary event whereby Bitcoin "goes to zero" with a normal distribution of the type we use. It is not clear why, after thirteen years, investors would suddenly decide that cryptocurrency is literally worthless. However, one can imagine events like the US government banning possession or trading of cryptocurrency, similar to how private possession of gold was greatly curtailed in the US in 1933.

The common intuition of the effect of cryptocurrency "going to zero" is that it will wipe out all the gains from a long position. However, this only applies for buy-and-hold positions, rather than a rebalanced portfolio. For instance, if one holds cryptocurrency at a weight of 1% and it doubles in price (before eventually going to zero), the investor responds to the higher portfolio weight by selling cryptocurrency and buying equities - in effect, locking in part of one's gains into the equity portfolio. The maximum downside exposure in a single month of a long position is only the chosen weight, which is generally small. For an investor who rebalances monthly, going to zero in a single month is much less problematic than declining by 90% per month for two years straight, where the rebalancing effect increases weights each time before further losses. In other words, the disaster scenarios of a rebalanced portfolio differ from those commonly described.

Events like this will both reduce mean returns and add significant non-normality. It is difficult to know the correlation of such an event with equity returns. A zero correlation is simple, but seems unrealistic. Even if the event itself occurred randomly (which was not true for gold restrictions), cryptocurrency is a large enough asset class that its value going to zero likely would have spillover financial effects. Similar arguments apply for other sources of non-normality, like disaster risk (Fagereng, Gottlieb, and Guiso 2017). In principle, these ought to be amenable to suitably modified Bayesian portfolio optimization tools, although the ease and feasibility of this is unclear.

Nevertheless, such alternatives come up against a powerful general intuition. It is easy to imagine disaster events that make *buying* cryptocurrency unattractive, but these ought to make *shorting* cryptocurrency *more attractive*. The most likely extensions that could justify zero weights are those that imply that the costs of *any* trade or position are larger than we think, or that the distribution of returns is more volatile in both directions.

Just as no single paper resolved the stock market participation puzzle, we do not address all possible reasons why investors may not trade cryptocurrency. The role of different forms of preferences is an interesting open question. In the stock market participation literature, these include loss aversion (F. J. Gomes 2005), narrow framing (Barberis, Huang, and Thaler 2006), ambiguity aversion (Epstein and T. Wang 1994), rank dependence (Chapman and Polkovnichenko 2009, disappointment aversion (Ang, Bekaert, and J. Liu 2005) and news utility (Pagel 2018).

We study the actions implied by various beliefs, but do not take a stand on *why* investors have the priors they do, nor if such priors are reasonable. Much of the public debate involves (ineffectually) trying convince others that their priors are wrong. While we try to avoid this, some readers may be reluctant to invest without a clear economic theory of the asset. In the Appendix we suggest an economic basis for at least agnosticism about cryptocurrency returns. However, none of our analysis depends on this, and we assume that everyone is entitled to their priors.

6 Conclusion

A puzzling pattern in household finance is that more than 76% of people in the U.S. do not invest in the cryptocurrency market. Individual cryptocurrency participation is much lower than would be predicted by models such as the Consumption Capital Asset Pricing Model (CCAPM), given the risk-adjusted expected returns from holding cryptocurrency assets. Insights into the causes of nonparticipation may guide efforts to more effectively promote efficient financial decision making.

The above sentences are not, in fact, our own. They are lifted directly from papers on the stock market participation puzzle, with the words "stock" or "equity" substituted with "cryptocurrency".[10] To many economists, the comparison will seem absurd. But *why* should it be absurd? It is not because the sentences are literally false, if interpreted straightforwardly as statements about historical asset returns. Rather, resistance seems to arise from intuitions about the economic nature of the asset classes. Public firms hold productive assets, produce ongoing cash flows, and drive a vast amount of real economic activity. In contrast, most cryptocurrencies exist solely as numbers on a computer, could be trivially forked to produce alternative versions, hold no physical assets, and produce no cash flows because they are incapable by design of doing so.

Lacking good models of cryptocurrency returns, much of mainstream finance has tended towards incredulity. Venerable asset pricing tools like discounted cashflows say the price should be zero, so the returns make no sense, and only naive or confused investors purchase them. This attitude has faced an uneasy tension, as investors in these purportedly nonsensical assets have outperformed over long horizons equity portfolios formed on traditional academic advice. If a student explains their beliefs about cryptocurrency and asks what their allocation should be, what answer should one give? Positive? Negative? Zero? Not sure? We argue that it is time to tackle this tension head-on.

We explore this question using another equally venerable set of asset pricing tools – portfolio theory. The lessons learned stand in sharp contrast to much of the popular discussion. Zero weights are more difficult to justify than many investors presume. They are especially difficult to combine with strong professed beliefs that cryptocurrency is a bubble. In particular, the refusal to take even small positions, long or short, is puzzling. Directionally, close-to-zero weights are most easily generated by either i) dogmatic beliefs that cryptocurrency will earn slightly less than the risk-free rate (a position not frequently publicly espoused), or ii) extremely high frictions and costs, greater than most plausible estimates. The alternative is that zero weights in cryptocurrencies represent a behavioral rule of thumb, rather than an optimal choice under realistic frictions.

[10] Respectively, Kuhnen and Miu 2017, Bogan 2008, and Briggs et al. 2021. Around 50% of households do not hold equities (Campbell 2006, Calvet, Campbell, and Sodini 2007).

References

Abel, Andrew B, Janice C Eberly, and Stavros Panageas (2013). "Optimal inattention to the stock market with information costs and transactions costs". In: *Econometrica* 81.4, pp. 1455–1481.

Abreu, Dilip and Markus K Brunnermeier (2003). "Bubbles and crashes". In: *Econometrica* 71.1, pp. 173–204.

Anderson, Evan W and Ai-Ru Cheng (2016). "Robust Bayesian portfolio choices". In: *The Review of Financial Studies* 29.5, pp. 1330–1375.

Anderson, Evan W, Lars Peter Hansen, and Thomas J Sargent (2003). "A quartet of semigroups for model specification, robustness, prices of risk, and model detection". In: *Journal of the European Economic Association* 1.1, pp. 68–123.

Ang, Andrew, Geert Bekaert, and Jun Liu (2005). "Why stocks may disappoint". In: *Journal of Financial Economics* 76.3, pp. 471–508.

Barberis, Nicholas, Ming Huang, and Richard H Thaler (2006). "Individual preferences, monetary gambles, and stock market participation: A case for narrow framing". In: *American economic review* 96.4, pp. 1069–1090.

Barberis, Nicholas, Andrei Shleifer, and Robert Vishny (1998). "A model of investor sentiment". In: *Journal of financial economics* 49.3, pp. 307–343.

Biais, Bruno et al. (2022). "Equilibrium bitcoin pricing". In: *Journal of Finance*.

Bogan, Vicki (2008). "Stock market participation and the internet". In: *Journal of Financial and Quantitative Analysis* 43.1, pp. 191–211.

Boiko, Viktor et al. (2021). "The optimization of the cryptocurrency portfolio in view of the risks". In: *Journal of Management Information and Decision Sciences. 2021, № 4, vol. 24. . 1-9.*

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer (2018). "Diagnostic expectations and credit cycles". In: *The Journal of Finance* 73.1, pp. 199–227.

Brauneis, Alexander and Roland Mestel (2019). "Cryptocurrency-portfolios in a mean-variance framework". In: *Finance Research Letters* 28, pp. 259–264.

Briggs, Joseph et al. (2021). "Windfall gains and stock market participation". In: *Journal of Financial Economics* 139.1, pp. 57–83.

Calvet, Laurent E, John Y Campbell, and Paolo Sodini (2007). "Down or out: Assessing the welfare costs of household investment mistakes". In: *Journal of Political Economy* 115.5, pp. 707–747.

Campbell, John Y (2006). "Household finance". In: *The journal of finance* 61.4, pp. 1553–1604.

Carhart, Mark M (1997). "On persistence in mutual fund performance". In: *The Journal of finance* 52.1, pp. 57–82.

Chapman, David A and Valery Polkovnichenko (2009). "First-order risk aversion, heterogeneity, and asset market outcomes". In: *The Journal of Finance* 64.4, pp. 1863–1887.

Chuen, David LEE Kuo, Li Guo, and Yu Wang (2017). "Cryptocurrency: A new investment opportunity?" In: *The journal of alternative investments* 20.3, pp. 16–40.

Cong, Lin William, Ye Li, and Neng Wang (2021). "Tokenomics: Dynamic adoption and valuation". In: *The Review of Financial Studies* 34.3, pp. 1105–1155.

Daniel, Kent, David Hirshleifer, and Avanidhar Subrahmanyam (1998). "Investor psychology and security market under-and overreactions". In: *the Journal of Finance* 53.6, pp. 1839–1885.

Epstein, Larry G and Tan Wang (1994). “Intertemporal Asset Pricing under Knightian Uncertainty”. In: *Econometrica* 62.2, pp. 283–322.
Fagereng, Andreas, Charles Gottlieb, and Luigi Guiso (2017). “Asset market participation and portfolio choice over the life-cycle”. In: *The Journal of Finance* 72.2, pp. 705–750.
Fama, Eugene F and Kenneth R French (1993). “Common risk factors in the returns on stocks and bonds”. In: *Journal of financial economics* 33.1, pp. 3–56.
Flori, Andrea (2019). “Cryptocurrencies in finance: Review and applications”. In: *International Journal of Theoretical and Applied Finance* 22.05, p. 1950020.
Foley, Sean, Jonathan R Karlsen, and Tālis J Putniņš (2019). “Sex, drugs, and bitcoin: How much illegal activity is financed through cryptocurrencies?” In: *The Review of Financial Studies* 32.5, pp. 1798–1853.
French, Kenneth R and James M Poterba (1991). “Investor Diversification and International Equity Markets”. In: *The American Economic Review* 81.2, pp. 222–226.
Gabaix, Xavier and Ralph SJ Koijen (2021). *In search of the origins of financial fluctuations: The inelastic markets hypothesis*. Tech. rep. National Bureau of Economic Research.
Gomes, Francisco and Alexander Michaelides (2005). “Optimal life-cycle asset allocation: Understanding the empirical evidence”. In: *The Journal of Finance* 60.2, pp. 869–904.
Gomes, Francisco J (2005). “Portfolio choice and trading volume with loss-averse investors”. In: *The Journal of Business* 78.2, pp. 675–706.
Greenwood, Robin, Andrei Shleifer, and Yang You (2019). “Bubbles for fama”. In: *Journal of Financial Economics* 131.1, pp. 20–43.
Griffin, John M and Amin Shams (2020). “Is Bitcoin really untethered?” In: *The Journal of Finance* 75.4, pp. 1913–1964.
Haliassos, Michael and Alexander Michaelides (2003). “Portfolio choice and liquidity constraints”. In: *International Economic Review* 44.1, pp. 143–177.
Hansen, Lars Peter and Thomas J Sargent (2001). “Robust control and model uncertainty”. In: *American Economic Review* 91.2, pp. 60–66.
— (2011). “Robustness”. In: *Robustness*. Princeton university press.
Hartzmark, Samuel M and David H Solomon (2021). “Predictable price pressure”. In: *Available at SSRN 3853096*.
Harvey, Campbell R et al. (2022). “An Investor’s Guide to Crypto”. In: *Available at SSRN 4124576*.
Hong, Harrison and David Sraer (2013). “Quiet bubbles”. In: *Journal of Financial Economics* 110.3, pp. 596–606.
Hong, Harrison and Jeremy C Stein (1999). “A unified theory of underreaction, momentum trading, and overreaction in asset markets”. In: *The Journal of finance* 54.6, pp. 2143–2184.
Hrytsiuk, Petro, Tetiana Babych, and Larysa Bachyshyna (2019). “Cryptocurrency portfolio optimization using Value-at-Risk measure”. In: *Advances in Economics, Business and Management Research* 95, pp. 385–389.
Kuhnen, Camelia M and Andrei C Miu (2017). “Socioeconomic status and learning from financial information”. In: *Journal of Financial Economics* 124.2, pp. 349–372.
Li, Jiasun and William Mann (2018). “Digital tokens and platform building”. In.

Liu, Yukun and Aleh Tsyvinski (2021). “Risks and returns of cryptocurrency”. In: *The Review of Financial Studies* 34.6, pp. 2689–2727.

Liu, Yukun, Aleh Tsyvinski, and Xi Wu (Apr. 2022). “Common Risk Factors in Cryptocurrency”. In: *The Journal of Finance* 77.2, pp. 1133–1177.

Makarov, Igor and Antoinette Schoar (2019). “Price discovery in cryptocurrency markets”. In: *AEA Papers and Proceedings*. Vol. 109, pp. 97–99.

— (2020). “Trading and arbitrage in cryptocurrency markets”. In: *Journal of Financial Economics* 135.2, pp. 293–319.

— (2021). *Blockchain analysis of the bitcoin market*. Tech. rep. National Bureau of Economic Research.

Mankiw, N Gregory and Stephen P Zeldes (1991). “The consumption of stockholders and non-stockholders”. In: *Journal of financial Economics* 29.1, pp. 97–112.

Markowitz, Harry (1952). “The utility of wealth”. In: *Journal of political Economy* 60.2, pp. 151–158.

Nakamoto, Satoshi (2008). “Bitcoin: A peer-to-peer electronic cash system”. In: *Decentralized Business Review*, p. 21260.

Pagel, Michaela (2018). “A News-Utility Theory for Inattention and Delegation in Portfolio Choice”. In: *Econometrica* 86.2, pp. 491–522.

Pagnotta, Emiliano S (2022). “Decentralizing money: Bitcoin prices and blockchain security”. In: *The Review of Financial Studies* 35.2, pp. 866–907.

Pástor, L'uboš (2000). “Portfolio selection and asset pricing models”. In: *The Journal of Finance* 55.1, pp. 179–223.

Petukhina, Alla et al. (2021). “Investing with cryptocurrencies–evaluating their potential for portfolio allocation strategies”. In: *Quantitative Finance* 21.11, pp. 1825–1853.

Rozario, Evans et al. (2020). “A Decade of Evidence of Trend Following Investing in Cryptocurrencies”. In: *Available at SSRN 3697981*.

Scheinkman, Jose A and Wei Xiong (2003). “Overconfidence and speculative bubbles”. In: *Journal of political Economy* 111.6, pp. 1183–1220.

Sockin, Michael and Wei Xiong (2020). “A model of cryptocurrencies”. In.

Tesar, Linda L and Ingrid M Werner (1995). “Home bias and high turnover”. In: *Journal of international Money and Finance* 14.4, pp. 467–492.

Tu, Jun and Guofu Zhou (2010). “Incorporating economic objectives into Bayesian priors: Portfolio choice under parameter uncertainty”. In: *Journal of Financial and Quantitative Analysis* 45.4, pp. 959–986.

Vissing-Jorgensen, Annette (2003). “Perspectives on behavioral finance: Does" irrationality" disappear with wealth? Evidence from expectations and actions”. In: *NBER macroeconomics annual* 18, pp. 139–194.

Yermack, David (2017). “Corporate governance and blockchains”. In: *Review of finance* 21.1, pp. 7–31.

Yi, Shuyue, Zishuang Xu, and Gang-Jin Wang (2018). “Volatility connectedness in the cryptocurrency market: Is Bitcoin a dominant cryptocurrency?” In: *International Review of Financial Analysis* 60, pp. 98–114.

Zhang, Wei et al. (2018). “Some stylized facts of the cryptocurrency market”. In: *Applied Economics* 50.55, pp. 5950–5965.

Figure 1: Prior Beliefs and Zero Investment in Cryptocurrency

This figure plots the time series of cutoff prior beliefs about the average monthly cryptocurrency return at the beginning of the sample period required for non-investment in cryptocurrency. Panel A plots the beliefs required for non-investment at each point in time, whereas Panel B plots the beliefs required for non-investment at any point up to each point in time. If the priors are above (below) the cutoff level, then investors should long (short) on a specific date (Panel A) or at some point prior to the date (Panel B). The calculations assume the following: (1) Investors start with the CRSP value-weighted market portfolio as a base asset and consider adding cryptocurrencies to their portfolios; (2) Investors observed ten years of data with a mean equal to their prior mean before the beginning of the sample period; (3) The variance of cryptocurrency returns approximately equals their ex-post variance – 150 times the market variance for Bitcoin, 170 times the market variance for the value-weighted cryptocurrency portfolio, and 625 times the market variance for the equal-weighted portfolio; (4) Investors believe cryptocurrency to be uncorrelated with the market portfolio. The sample consists of 106 monthly returns from May 2013 to February 2022.

Panel A: Snapshot Non-Investment

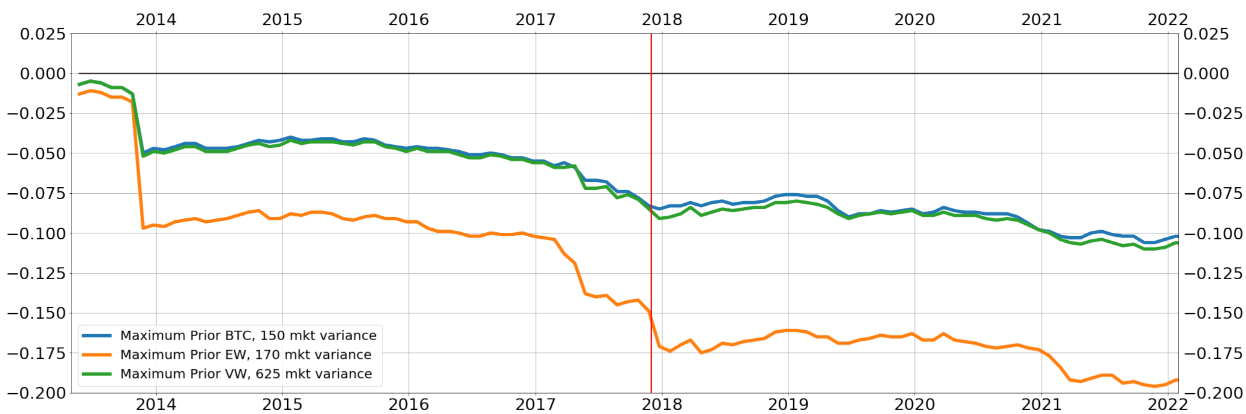

Panel B: Cumulative Non-Investment

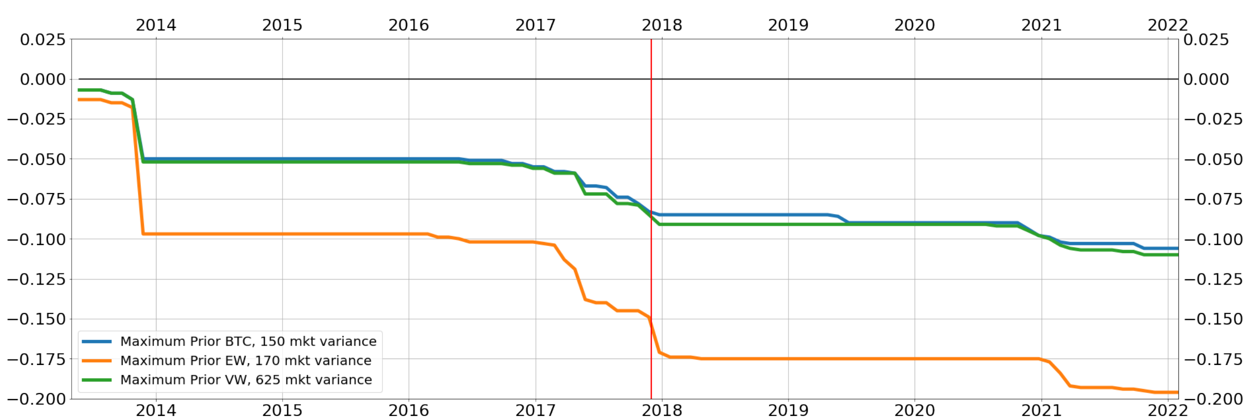

Figure 2: Optimal Cryptocurrency Portfolio Weights

This figure plots the time series of optimal cryptocurrency portfolio weights for different prior beliefs about the average monthly cryptocurrency return at the beginning of the sample period. Panel A shows the optimal weights for Bitcoin, whereas Panel B shows the optimal absolute weights for the equally-weighted cryptocurrency portfolio. Panels C and D show how the maximum desired weights over the sample vary continuously with priors. Panels E and F show the same thing with short sales constraints. Optimal weights are calculated for prior means between 2% per month and -20% per month, with a strength equal to ten years of prior data.

Panel A: Bitcoin

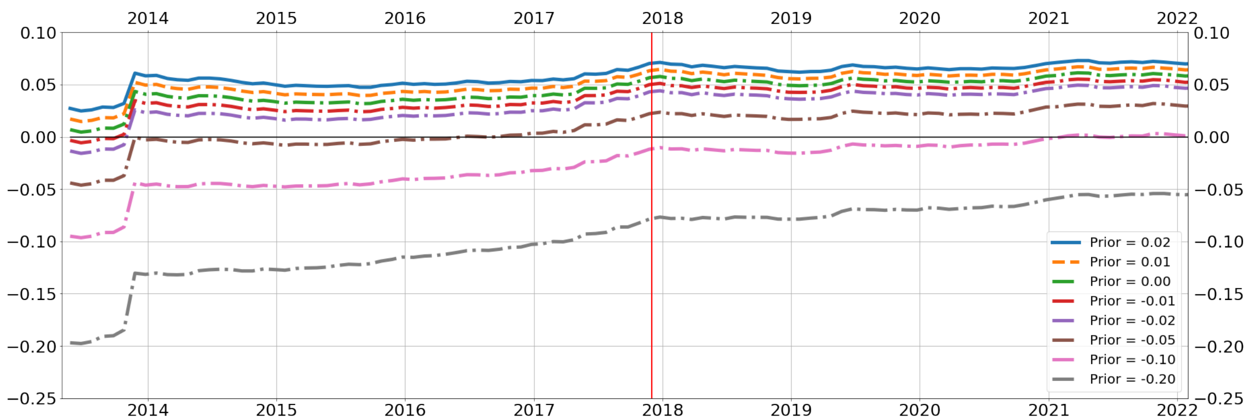

Panel B: Equal-weighted Cryptocurrency Portfolio

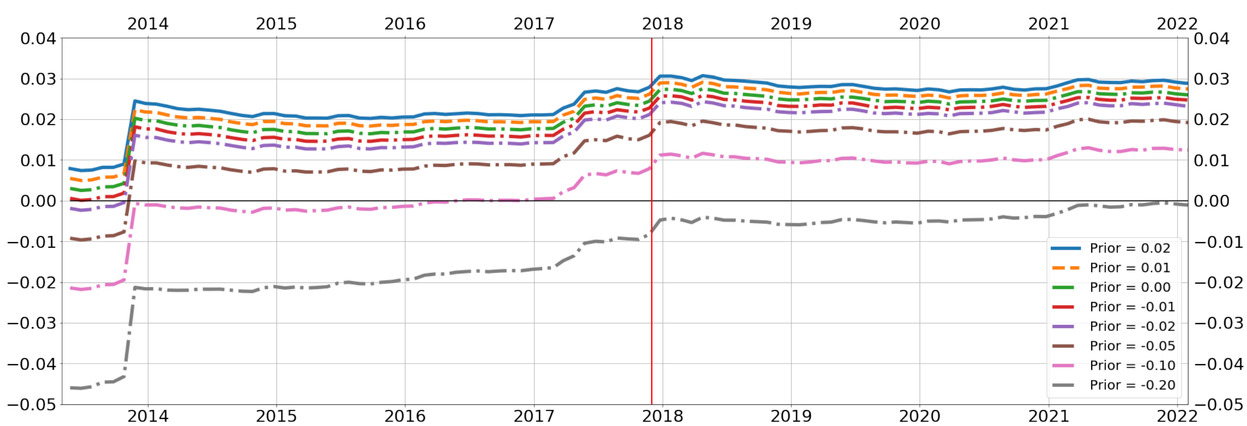

Panel C: Maximum absolute weights of BTC

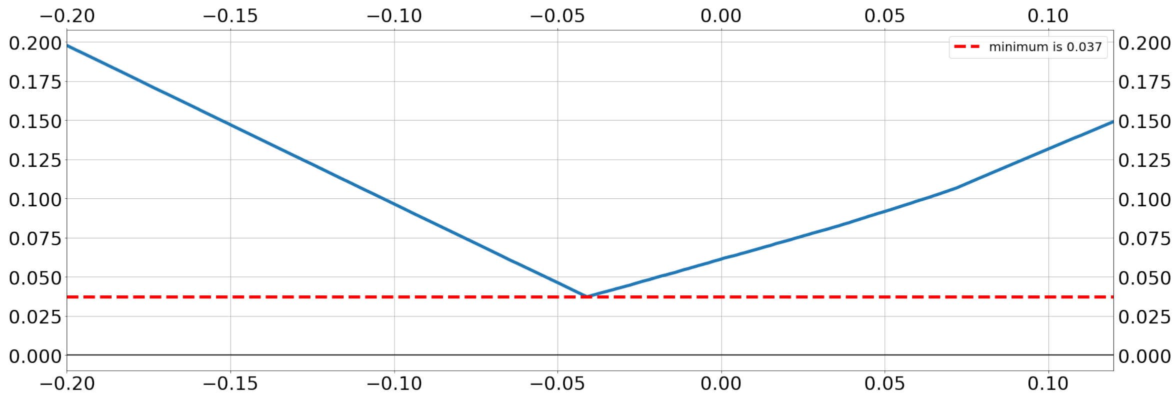

Panel D: Maximum absolute weights of Equal-weighted Cryptocurrency Portfolio

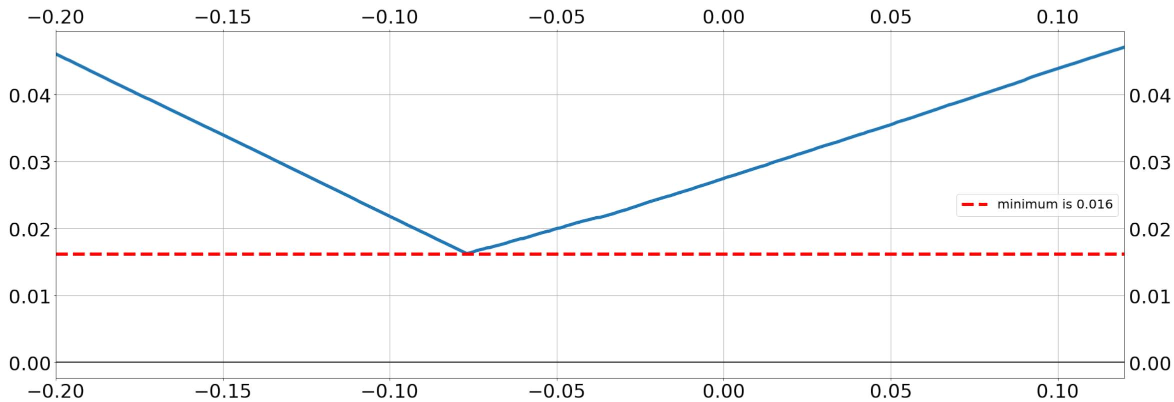

Panel E: Maximum absolute weights of BTC, no short before 2017

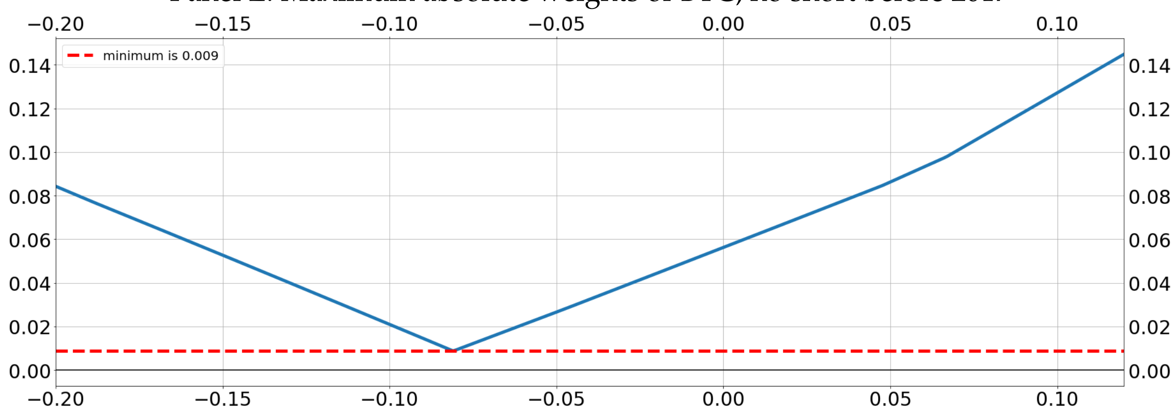

Panel F: Maximum absolute weights of Equal-weighted Cryptocurrency, no short

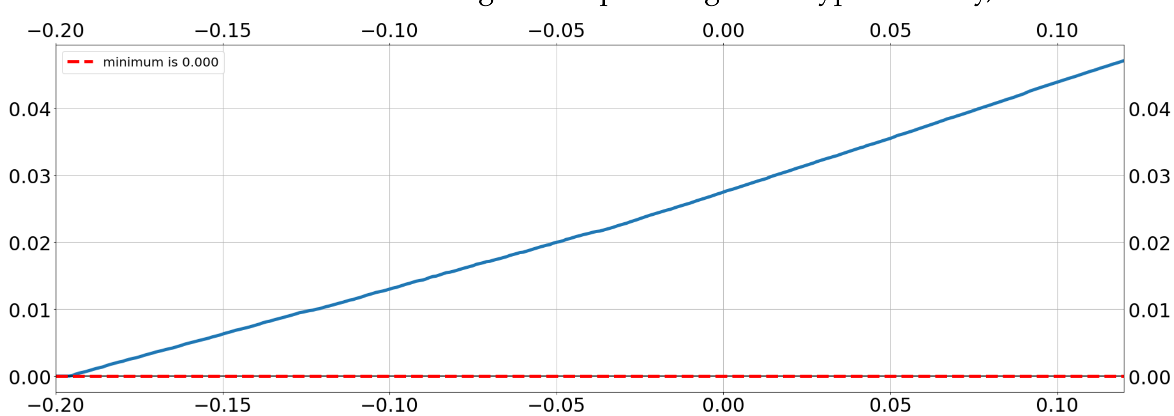

Figure 3: Ex-Ante Certainty Equivalent Gains from Cryptocurrency

This figure plots the end-of-sample certainty equivalent of return (CER) gains from adding cryptocurrencies to investors' existing portfolios, as a function of different prior beliefs about the average monthly cryptocurrency return at the beginning of the sample period. The reported values equal the difference between the CER of the baseline market portfolio that excludes cryptocurrency and the CER of the optimal portfolio that combines the market portfolio and cryptocurrency, assuming that investors have a constant relative risk aversion of 3. Panel A shows the CER gains for Bitcoin, whereas Panel B shows the CER gains for the equally-weighted cryptocurrency portfolio.

Panel A:Bitcoin

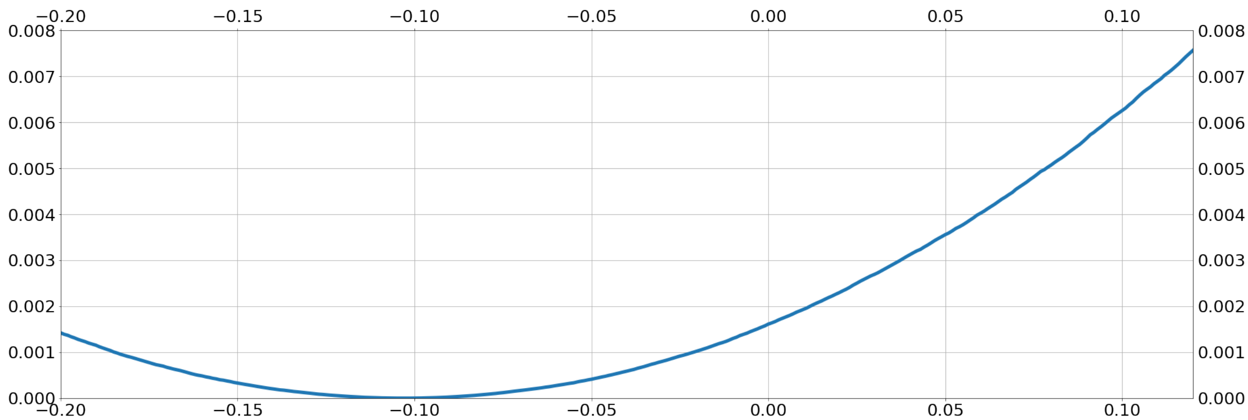

Panel B: Equal-weighted Cryptocurrency Portfolio

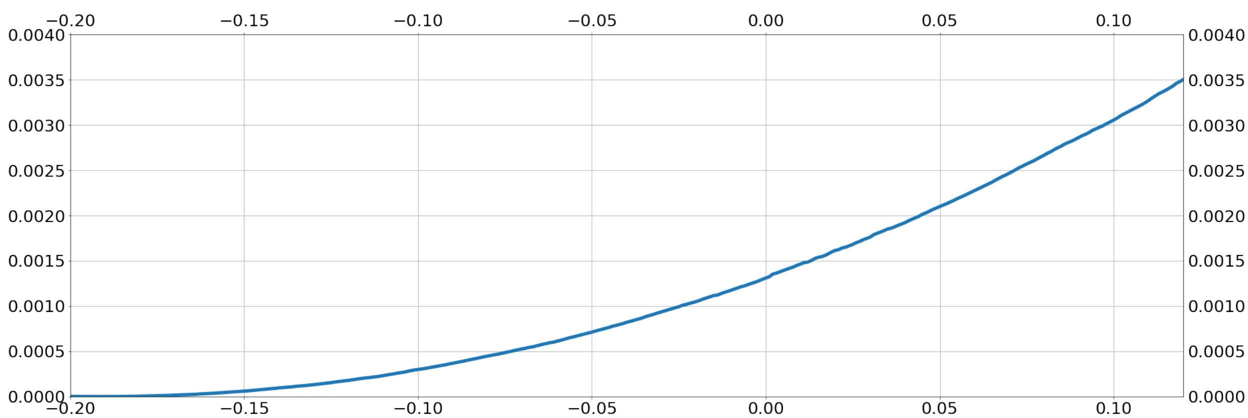

Figure 4: Ex-Post Certainty Equivalent Gains from Cryptocurrency

This figure plots the end-of-sample ex-post certainty equivalent return (CER) gains from adding cryptocurrencies to investors' existing portfolios across different prior beliefs about the average monthly cryptocurrency return at the beginning of the sample period. Investors assess ex-post performance on a distributional basis, assuming that the distribution of realized returns up to that point (from whatever series of weights was chosen) were to continue indefinitely. The reported values equal the difference between the ex-post CER of the baseline market portfolio that excludes cryptocurrency and the ex-post CER of the optimal portfolio that combines the market portfolio and cryptocurrency. Investors are assumed to have a constant relative risk aversion of 3. Panel A shows the ex-post CER gains for Bitcoin, whereas Panel B shows the ex-post CER gains for the equally-weighted cryptocurrency portfolio.

Panel A:Bitcoin

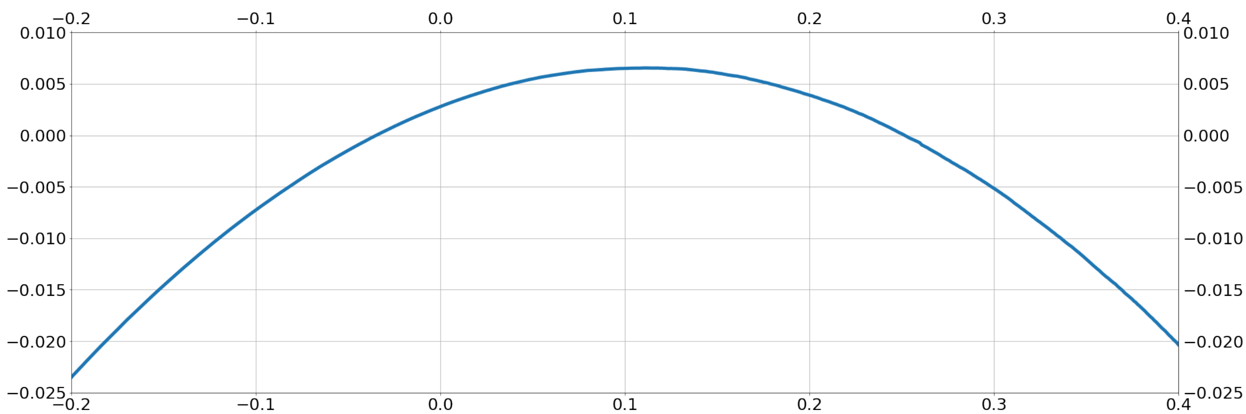

Panel B: Equal-weighted Cryptocurrency Portfolio

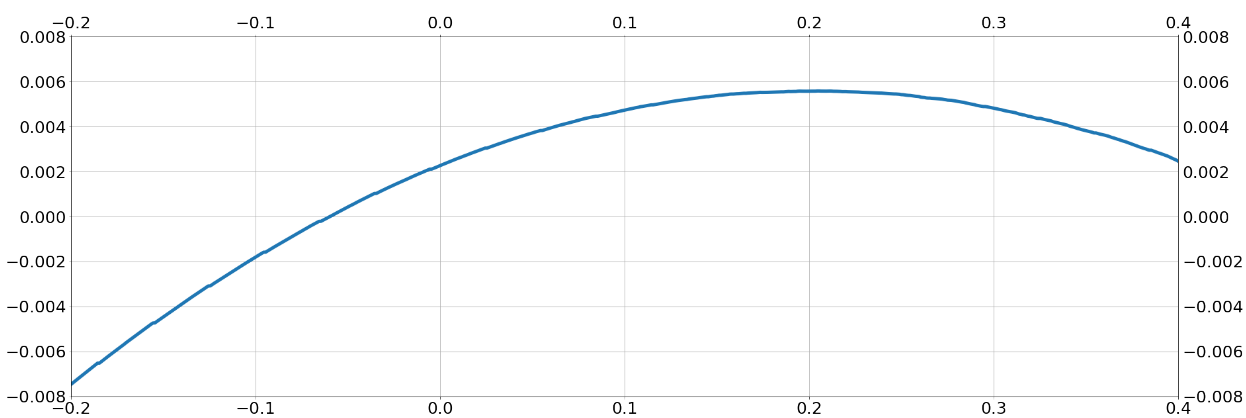

Figure 5: Ex-ante vs. Ex-post Performance of Cryptocurrency

This figure plots together ex-ante and ex-post assessments of the performance of cryptocurrency portfolios. For each date, the figure plots the range of prior beliefs for which investors were: (1) long or short cryptocurrencies, on average, prior to that date; (2) long or short cryptocurrencies on that date; and (3) happy about their investment decision. Panel A focuses on Bitcoin, whereas Panel B focuses on the equally-weighted cryptocurrency portfolio. Green regions are when the investor was long on average up to that point, and also long at that point. Yellow regions are when the investor was short on average up to that point, but long at that point. Red regions are when the investor was short on average up to that point, and short at that point. Meanwhile, shaded regions are where the investor was ex-post happy with their distribution of returns up to that point, and unshaded regions are where the investor was unhappy ex-post up to that point.

Panel A:Bitcoin

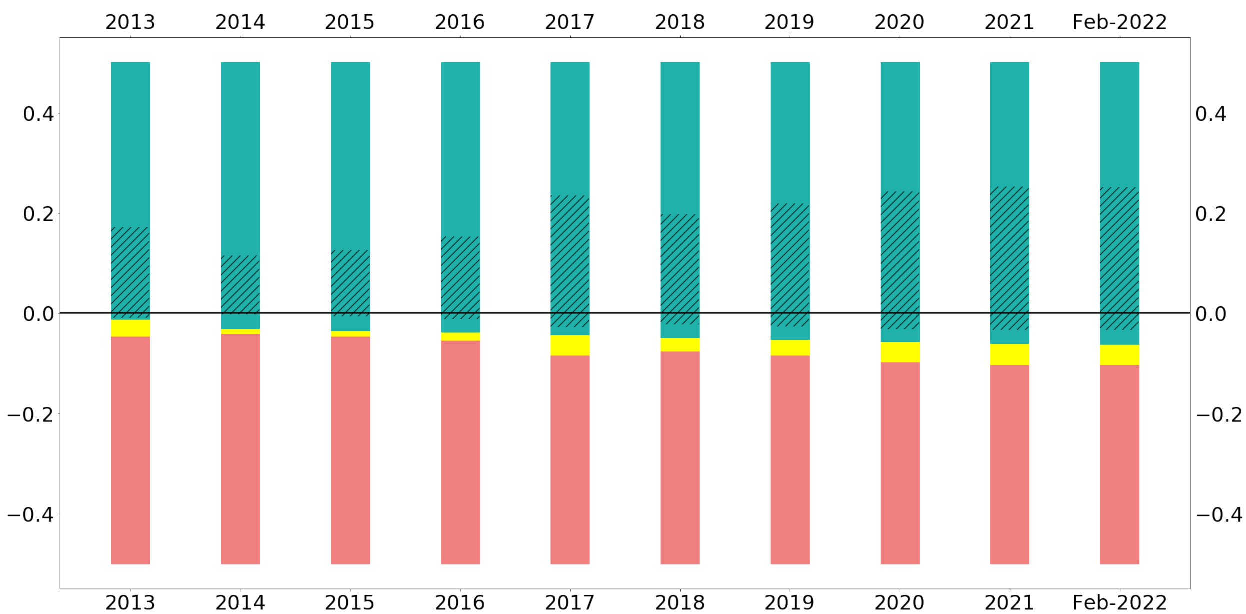

Panel B: Equal-weighted Cryptocurrency Portfolio

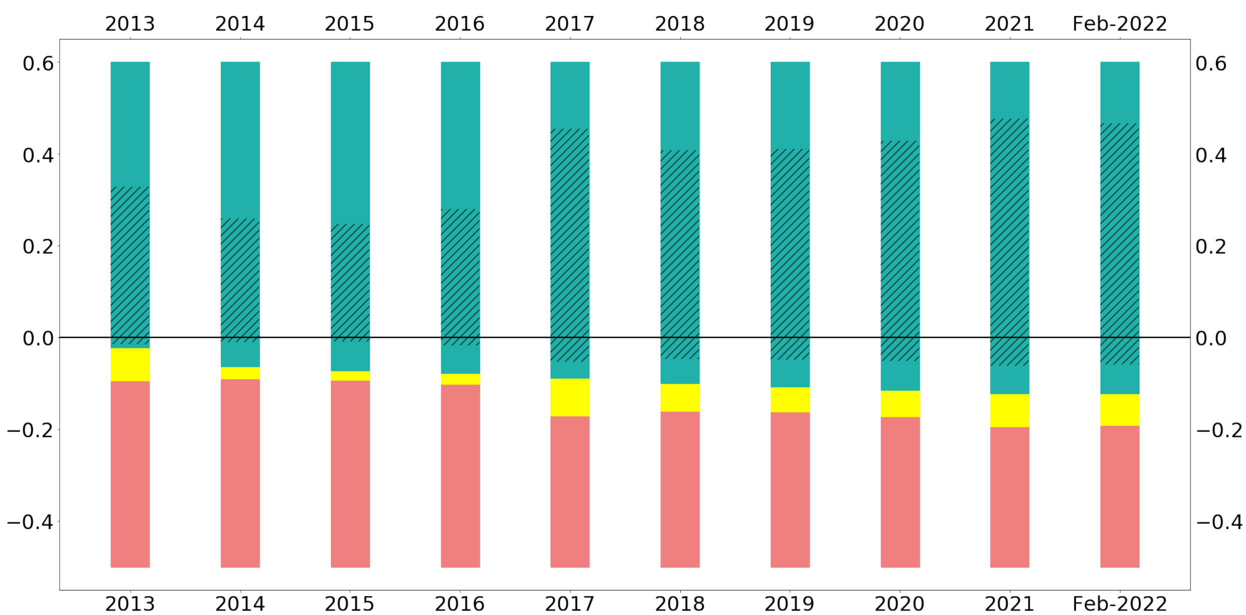

Table 1: Summary Statistics

This table reports summary statistics (Panel A) and correlations (Panel B) for monthly excess returns on the following portfolios:Bitcoin (BTC-RF), Equally-weighted cryptocurrencies (ew100-rf), Value-weighted cryptocurrencies (vw100-rf), CRSP value-weighted equity portfolio (mkt-rf), Small minus big portfolio (SMB), High minus low portfolio (HML), Momentum portfolio (UMD), and MSCI world ex-US portfolio (MSCI-rf). The sample consists of 106 monthly returns from May 2013 to February 2022.

Panel A: Descriptive Statistics

	BTC-RF	ew100-rf	vw100-rf	mkt-rf	smb	hml	umd	MSCI-rf
count	106	106	106	106	106	106	106	106
mean	0.110	0.201	0.113	0.011	0.000	-0.002	0.001	0.002
std	0.494	1.020	0.532	0.041	0.026	0.034	0.037	0.041
min	-0.382	-0.462	-0.407	-0.134	-0.059	-0.139	-0.124	-0.147
25%	-0.091	-0.145	-0.112	-0.009	-0.020	-0.020	-0.021	-0.021
50%	0.035	0.039	0.029	0.014	0.003	-0.005	0.004	0.007
75%	0.234	0.214	0.211	0.032	0.018	0.014	0.022	0.029
max	4.493	9.582	4.735	0.137	0.071	0.127	0.100	0.152
SR	0.224	0.198	0.214	0.279	-0.007	-0.052	0.039	0.061
skew	6.719	7.601	6.425	-0.358	0.215	0.268	-0.287	-0.150
kurtosis	56.924	66.139	52.367	1.722	-0.329	3.344	0.970	2.051

Panel B: Correlations

	BTC-RF	ew100-rf	vw100-rf	mkt-rf	smb	hml	umd	MSCI-rf
BTC-RF	1.000	0.915	0.978	0.165	0.031	-0.014	0.003	0.133
ew100-rf	0.915	1.000	0.953	0.108	0.026	-0.014	-0.013	0.094
vw100-rf	0.978	0.953	1.000	0.164	0.016	-0.021	-0.001	0.143
mkt-rf	0.165	0.108	0.164	1.000	0.313	0.060	-0.372	0.871
smb	0.031	0.026	0.016	0.313	1.000	0.073	-0.214	0.213
hml	-0.014	-0.014	-0.021	0.060	0.073	1.000	-0.509	0.130
umd	0.003	-0.013	-0.001	-0.372	-0.214	-0.509	1.000	-0.447
MSCI-rf	0.133	0.094	0.143	0.871	0.213	0.130	-0.447	1.000

Table 2: Prior Beliefs Leading to Zero Investment in Cryptocurrency

This table reports cutoff prior beliefs about the average monthly cryptocurrency return at the beginning of the sample period that would make an investor not invest in cryptocurrencies. Each row corresponds to a different cryptocurrency portfolio (Bitcoin, Equally-weighted cryptocurrency portfolio, Value-weighted cryptocurrency portfolio), and each column corresponds to a specific end-of-year date when the investment decision is made. Panel A calculates the cutoff prior belief that leads to no cryptocurrency investment on a specific date, whereas Panel B calculates the cutoff belief that leads to no cryptocurrency investment at any point prior to the date. If the priors are above (below) the cutoff level, then investors should long (short) on a specific date (Panel A) or at some point prior to the date (Panel B). The calculations assume the following: (1) Investors start with the CRSP value-weighted market portfolio as a base asset and consider adding cryptocurrencies to their portfolios; (2) Investors observed ten years of data with a mean equal to their prior mean before the beginning of the sample period; (3) The variance of cryptocurrency returns approximately equals their ex-post variance – 150 times the market variance for Bitcoin, 170 times the market variance for the value-weighted cryptocurrency portfolio, and 625 times the market variance for the equal-weighted portfolio; (4) Investors believe cryptocurrency to be uncorrelated with the market portfolio. The sample consists of 106 monthly returns from May 2013 to February 2022.

Panel A: Snapshot Non-Investment

	2013	2014	2015	2016	2017	2018	2019	2020	2021	End of Sample
Btc-rf	-0.047	-0.042	-0.047	-0.055	-0.085	-0.076	-0.085	-0.098	-0.104	-0.103
ew-rf	-0.095	-0.091	-0.093	-0.102	-0.171	-0.161	-0.163	-0.173	-0.195	-0.192
vw-rf	-0.049	-0.045	-0.049	-0.056	-0.091	-0.081	-0.086	-0.098	-0.109	-0.107

Panel B: Cumulative Non-Investment

	2013	2014	2015	2016	2017	2018	2019	2020	2021	End of Sample
Btc-rf	-0.05	-0.05	-0.05	-0.055	-0.085	-0.085	-0.09	-0.098	-0.106	-0.106
ew-rf	-0.097	-0.097	-0.097	-0.102	-0.171	-0.175	-0.175	-0.175	-0.196	-0.196
vw-rf	-0.052	-0.052	-0.052	-0.056	-0.091	-0.091	-0.091	-0.098	-0.110	-0.110

Table 3: Optimal Cryptocurrency Portfolio Weights

This table reports cryptocurrency portfolio weights for different prior beliefs about the average monthly cryptocurrency excess return at the beginning of the sample period. Panel A considers Bitcoin and Panel B considers an equally-weighted cryptocurrency portfolio. Each row corresponds to a different prior belief. For each prior, the columns indicate a range of attributes of the distribution of weights over the sample period - the average, lowest, highest, final (i.e. end of sample) weight, the fraction of months that are above zero, the fraction of weights whose absolute value exceeds 0.5%, 1.5%, 2.5% and 5%, the first date in the sample when the weight is positive, as well as the mean and standard deviation of leverage choices. The sample consists of 106 monthly returns from May 2013 to February 2022.

Panel A: Bitcoin

Prior	Average	S.t.d	Lowest	Highest	Final	Fraction positive	Fraction above 0.5%	Fraction above 1%	Fraction above 2%	Fraction above 5%	First Date Weight is Positive	Mean of Leverage	S.t.d of Leverage
Flat	0.160	0.090	-0.042	0.637	0.121	0.991	1.000	0.991	0.991	0.981	2013-06	1.085	0.097
2	0.059	0.011	0.025	0.073	0.070	1.000	1.000	1.000	1.000	0.830	all above	0.991	0.036
1	0.052	0.012	0.015	0.067	0.064	1.000	1.000	1.000	0.953	0.566	all above	0.982	0.037
0	0.045	0.013	0.004	0.061	0.059	1.000	0.991	0.953	0.943	0.481	all above	0.974	0.042
-1	0.038	0.014	-0.006	0.055	0.053	0.953	0.953	0.943	0.943	0.170	2013-10	0.965	0.045
-2	0.030	0.015	-0.016	0.049	0.047	0.943	1.000	0.991	0.792	0.000	2013-11	0.956	0.047
-5	0.009	0.018	-0.046	0.032	0.030	0.642	0.764	0.604	0.462	0.000	2016-06	0.943	0.058
-10	-0.026	0.024	-0.096	0.003	0.002	0.104	0.849	0.679	0.481	0.057	2021-02	0.906	0.063
-20	-0.096	0.035	-0.198	-0.054	-0.055	0.000	1.000	1.000	1.000	1.000	all below	0.838	0.069

Panel B: Equal-weighted Cryptocurrency

Prior	Average	S.t.d	Lowest	Highest	Final	Fraction positive	Fraction above 0.5%	Fraction above 1%	Fraction above 2%	Fraction above 5%	First Date Weight is Positive	Mean of Leverage	S.t.d of Leverage
Flat	0.075	0.044	-0.010	0.296	0.053	0.991	0.991	0.981	0.953	0.943	2013-06	1.009	0.051
2	0.024	0.005	0.007	0.031	0.029	1.000	1.000	0.943	0.943	0.000	all above	0.968	0.030
1	0.023	0.006	0.005	0.029	0.027	1.000	0.991	0.943	0.642	0.000	all above	0.963	0.029
0	0.021	0.006	0.003	0.027	0.026	1.000	0.943	0.943	0.566	0.000	all above	0.962	0.029
-1	0.019	0.006	0.000	0.026	0.025	1.000	0.943	0.943	0.547	0.000	all above	0.957	0.028
-2	0.018	0.006	-0.002	0.024	0.023	0.943	0.943	0.943	0.528	0.000	2013-11	0.955	0.028
-5	0.013	0.007	-0.010	0.020	0.019	0.943	1.000	0.566	0.000	0.000	2013-11	0.948	0.031
-10	0.004	0.008	-0.022	0.013	0.012	0.632	0.604	0.368	0.047	0.000	2016-06	0.939	0.032
-20	-0.013	0.011	-0.046	-0.001	-0.001	0.000	0.679	0.472	0.283	0.000	all below	0.936	0.037

Table 4: Ex-Ante Gains in Certainty Equivalent of Returns from Access to Cryptocurrency

This table reports the monthly certainty equivalent return (CER) gains, in percentage points, from adding cryptocurrency to investors' existing portfolios. Panel A is for Bitcoin, Panel B is for an equal-weighted portfolio of cryptocurrency. The reported values equal the difference between the CER of the baseline market portfolio that excludes cryptocurrency and the CER of the optimal portfolio that combines the market portfolio and cryptocurrency, assuming that investors have a constant relative risk aversion of 3. Years correspond to the end of the calendar year in question. Numbers in green map to positive portfolio weights, and numbers in black map to short weights. Each row corresponds to a different prior belief, and each column corresponds to a specific end-of-year date when the investment decision is made. The sample consists of 106 monthly returns from May 2013 to February 2022.

Panel A: Ex-Ante CER Gains for Bitcoin

Prior	2013	2014	2015	2016	2017	2018	2019	2020	2021	End of Sample
Flat	16.629	2.225	1.184	0.916	1.424	0.807	0.735	0.800	0.713	0.680
2	0.183	0.132	0.136	0.147	0.253	0.191	0.203	0.233	0.232	0.229
1	0.132	0.092	0.097	0.110	0.208	0.152	0.166	0.194	0.197	0.193
0	0.089	0.060	0.067	0.079	0.167	0.119	0.132	0.160	0.163	0.161
-1	0.055	0.034	0.041	0.053	0.131	0.089	0.104	0.129	0.133	0.130
-2	0.029	0.016	0.022	0.031	0.097	0.065	0.077	0.101	0.104	0.103
-5	0.000	0.003	0.000	0.001	0.028	0.014	0.022	0.039	0.043	0.041
-10	0.116	0.118	0.084	0.053	0.005	0.012	0.004	0.000	0.000	0.000
-20	0.944	0.860	0.693	0.544	0.301	0.311	0.242	0.177	0.145	0.142

Panel B: Ex-Ante CER Gains for Equal-weighted Cryptocurrency

Prior	2013	2014	2015	2016	2017	2018	2019	2020	2021	End of Sample
Flat	16.205	2.525	1.089	0.736	1.372	0.851	0.637	0.568	0.592	0.553
2	0.129	0.101	0.091	0.094	0.197	0.161	0.148	0.150	0.166	0.161
1	0.107	0.084	0.076	0.080	0.176	0.143	0.133	0.136	0.150	0.146
0	0.088	0.068	0.062	0.066	0.157	0.127	0.117	0.121	0.135	0.131
-1	0.070	0.054	0.049	0.054	0.138	0.111	0.103	0.106	0.121	0.117
-2	0.055	0.041	0.037	0.043	0.122	0.097	0.091	0.094	0.109	0.105
-5	0.020	0.014	0.013	0.017	0.078	0.060	0.056	0.060	0.074	0.071
-10	0.000	0.001	0.000	0.000	0.026	0.018	0.017	0.020	0.031	0.030
-20	0.107	0.099	0.082	0.061	0.005	0.007	0.006	0.003	0.000	0.000

Table 5: Comparison of Ex-Ante Certainty Equivalent Gains Across Portfolios

This table compares the monthly certainty equivalent return (CER) gains from adding each of the following portfolios to investors' existing market portfolios: first,Bitcoin (BTC), Equally-weighted cryptocurrency (EW), and Value-weighted cryptocurrency (VW), each for a range of different ten-year priors. These are compared to other portfolios - Small minus big portfolio (SMB), High minus low portfolio (HML), Momentum portfolio (UMD), MSCI world ex-US portfolio (MSCI), as well as combinations of the different equity portfolios. For MSCI, the reported number is the gain in CER. For every other portfolio, the number reported is the ratio of the CER gain under that portfolio to the CER gain of MSCI world (so MSCI world is treated as a baseline denominator of CER gains). The reported values equal the end-of-sample difference between the CER of the baseline market portfolio that excludes each asset and the CER of the optimal portfolio that combines the market portfolio and that asset, assuming that investors have a constant relative risk aversion of 3. To calculate the CER gains of the MSCI, SMB, HML, and UMD portfolios, we assume that investors start with flat priors and form their prior beliefs based on realized returns from when they became available and through the end of the sample period. In columns 2, 3, and 4, which correspond to the cryptocurrency portfolios, each row corresponds to a different prior belief about average cryptocurrency monthly returns. The sample consists of 106 monthly returns from May 2013 to February 2022.

Times(except MSCI)		BTC		EW		VW	
MSCI Gain(percentage)	0.104	BTC prior 0.020 Gain	2.195	EW prior 0.020 Gain	1.545	VW prior 0.020 Gain	2.021
SMB Gain	0.138	BTC prior 0.010 Gain	1.85	EW prior 0.010 Gain	1.401	VW prior 0.010 Gain	1.708
HML Gain	3.973	BTC prior 0.000 Gain	1.541	EW prior 0.000 Gain	1.253	VW prior 0.000 Gain	1.425
UMD Gain	4.397	BTC prior -0.010 Gain	1.243	EW prior -0.010 Gain	1.124	VW prior -0.010 Gain	1.165
SMB+HML Gain	4.371	BTC prior -0.020 Gain	0.983	EW prior -0.020 Gain	1.004	VW prior -0.020 Gain	0.942
SMB+HML+UMD Gain	11.744	BTC prior -0.050 Gain	0.393	EW prior -0.050 Gain	0.681	VW prior -0.050 Gain	0.406
SMB+HML+MSCI Gain	6.031	BTC prior -0.100 Gain	0.001	EW prior -0.100 Gain	0.286	VW prior -0.100 Gain	0.006
SMB+HML+UMD+MSCI Gain	12.917	BTC prior -0.200 Gain	1.361	EW prior -0.200 Gain	0.002	VW prior -0.200 Gain	1.133
		BTC flat prior Gain	6.507	EW flat prior Gain	5.289	VW flat prior Gain	6.069

Table 6: Investment Costs and Beliefs Required for Non-Investment

This table considers the effect of annual investment costs on investors' optimal investment in cryptocurrency portfolios. Panels A and B report the range of prior means that would justify non-investment in cryptocurrencies throughout the sample for different investment costs (in percent per year, applied to the absolute value of the weight in cryptocurrency). Rows consider costs ranging from 10% per year to 50% per year. "Lowest Preventing Cost" is the smallest cost for which there is some ten-year prior that would result in non-investment over the whole sample up to that point, with "Corresponding Prior" being the beliefs that map to this non-investment. Panels C and D report the inverse - for each of the priors in question, how high would costs have to be to result in non-investment up to that point? For Panels C and D, we assume that investors cannot short Bitcoin before December 2017, and cannot short equal-weighted cryptocurrency at all.

Panel A:Bitcoin

Percentage Cost/Time	2013	2014	2015	2016	2017	2018	2019	2020	2021	End of Sample
10 (min)	nan	nan	nan	nan	nan	nan	nan	nan	nan	nan
(max)	nan	nan	nan	nan	nan	nan	nan	nan	nan	nan
15 (min)	-2.90	-2.90	-2.90	-2.90	nan	nan	nan	nan	nan	nan
(max)	-2.40	-2.40	-2.40	-2.40	nan	nan	nan	nan	nan	nan
20 (min)	-3.80	-3.80	-3.80	-3.80	-3.80	-3.80	-3.80	nan	nan	nan
(max)	-1.50	-1.50	-1.50	-1.50	-3.60	-3.60	-3.60	nan	nan	nan
30 (min)	-5.50	-5.50	-5.50	-5.50	-5.50	-5.50	-5.50	-5.50	-5.50	-5.50
(max)	0.30	0.30	0.30	0.30	-1.20	-1.20	-1.20	-1.20	-1.40	-1.40
50 (min)	-8.90	-8.90	-8.90	-8.90	-8.90	-8.90	-8.90	-8.90	-8.90	-8.90
(max)	3.80	3.80	3.80	3.80	3.70	3.70	3.70	3.70	3.70	3.70
Lowest Preventing cost	13.3	13.3	13.3	13.3	19.4	19.4	19.4	20.1	21.3	21.3
Corresponding Prior	-2.6	-2.6	-2.6	-2.6	-3.7	-3.7	-3.7	-3.8	-4.0	-4.0

Panel B: Equal-Weighted Cryptocurrency

Percentage Cost/Time	2013	2014	2015	2016	2017	2018	2019	2020	2021	End of Sample
10 (min)	nan	nan	nan	nan	nan	nan	nan	nan	nan	nan
(max)	nan	nan	nan	nan	nan	nan	nan	nan	nan	nan
15 (min)	nan	nan	nan	nan	nan	nan	nan	nan	nan	nan
(max)	nan	nan	nan	nan	nan	nan	nan	nan	nan	nan
20 (min)	nan	nan	nan	nan	nan	nan	nan	nan	nan	nan
(max)	nan	nan	nan	nan	nan	nan	nan	nan	nan	nan
30 (min)	-6.1	-6.1	-6.1	-6.1	nan	nan	nan	nan	nan	nan
(max)	-4.5	-4.5	-4.5	-4.5	nan	nan	nan	nan	nan	nan
50 (min)	-9.5	-9.5	-9.5	-9.5	-9.5	-9.5	-9.5	-9.5	-9.5	-9.5
(max)	-0.9	-0.9	-0.9	-0.9	-4.7	-4.8	-4.8	-4.8	-4.8	-4.8
Lowest Preventing cost	25.2	25.2	25.2	25.2	38.5	38.9	38.9	38.9	38.9	38.9
Corresponding Prior	-5.3	-5.3	-5.3	-5.3	-7.5	-7.6	-7.6	-7.6	-7.6	-7.6

Panel C: Minimum Cost for Non-Investment in Bitcoin (in Percent)

Priors in Percentage	2013	2014	2015	2016	2017	2018	2019	2020	2021	End of Sample
2	39.39	39.39	39.39	39.39	42.66	42.66	42.66	42.66	42.66	42.66
1	33.7	33.7	33.7	33.7	38.69	38.69	38.69	38.69	38.69	38.69
0	28.11	28.11	28.11	28.11	34.65	34.65	34.65	34.65	34.76	34.76
-1	22.42	22.42	22.42	22.42	30.68	30.68	30.68	30.68	31.33	31.33
-2	16.72	16.72	16.72	16.72	26.47	26.47	26.47	26.47	28.10	28.10
-5	0	0	0	2.07	14.17	14.17	14.43	16.29	17.97	17.97
-10	0	0	0	0	0	9.26	9.32	9.32	9.32	9.32
-20	0	0	0	0	0	48.22	48.22	48.22	48.22	48.22

Panel D: Minimum Cost for Non-Investment in Equal-Weighted Cryptocurrency

Priors in Percentage	2013	2014	2015	2016	2017	2018	2019	2020	2021	End Sample
2	66.59	66.59	66.59	66.59	77.20	77.20	77.20	77.20	77.20	77.20
1	60.77	60.77	60.77	60.77	72.96	73.07	73.07	73.07	73.07	73.07
0	55.05	55.05	55.05	55.05	68.90	69.01	69.01	69.01	69.01	69.01
-1	49.35	49.35	49.35	49.35	64.80	65.00	65.00	65.00	65.00	65.00
-2	43.75	43.75	43.75	43.75	60.76	61.10	61.10	61.10	61.10	61.10
-5	26.61	26.61	26.61	26.61	48.54	49.16	49.16	49.16	49.16	49.16
-10	0.00	0.00	0.00	0.77	28.35	29.32	29.32	29.32	30.91	30.91
-20	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Table 7: Ex-Post Certainty Equivalent Gains

This table reports investors' certainty equivalent return (CER) gains for the ex-post distribution of portfolio returns relative to the ex-post distribution of the market portfolio over the same period. Panel A examines Bitcoin, and Panel B examines equal-weighted cryptocurrency. The reported values equal the CER gains across annual snapshots over the sample period for different priors. Intuitively, the calculations assume that the distribution of portfolio returns that investors received up to each time point would continue indefinitely, and assess the CER gain over the distribution to that of the equity market portfolio alone. "Max Gain" calculates the maximum possible ex-post gain and "Optimal Prior" is the prior beliefs that correspond to this maximum. "Positive CER, max prior, T0=10" and the equivalent "min prior" describe the range of priors that map to ex-post positive CER gains at that point in time. Investors are assumed to have a relative risk aversion of 3. The sample consists of 106 monthly returns from May 2013 to February 2022.

Panel A: Bitcoin

Prior	2013	2014	2015	2016	2017	2018	2019	2020	2021	End of Sample
Flat	-2.212	-2.655	-1.489	-0.805	0.157	-0.156	0.014	0.166	0.213	0.199
2	1.600	0.347	0.275	0.306	0.594	0.370	0.391	0.429	0.418	0.405
1	1.126	0.221	0.189	0.229	0.494	0.301	0.324	0.362	0.356	0.344
0	0.574	0.062	0.083	0.136	0.382	0.221	0.249	0.288	0.287	0.278
-1	-0.046	-0.128	-0.041	0.031	0.260	0.133	0.166	0.208	0.213	0.205
-2	-0.741	-0.349	-0.186	-0.089	0.125	0.035	0.075	0.120	0.132	0.127
-5	-3.249	-1.193	-0.734	-0.531	-0.346	-0.316	-0.248	-0.186	-0.149	-0.148
-10	-8.847	-3.207	-2.028	-1.546	-1.354	-1.090	-0.948	-0.836	-0.742	-0.727
-20	-25.423	-9.541	-6.057	-4.624	-4.210	-3.349	-2.960	-2.670	-2.404	-2.353
Max Gain	2.962	0.556	0.428	0.483	0.992	0.588	0.627	0.693	0.677	0.654
Optimal Prior	8.1	5.6	6.0	6.9	10.4	8.4	9.3	10.5	11.1	11.2
CER>0, max prior	0.171	0.115	0.125	0.152	0.235	0.197	0.219	0.243	0.253	0.251
CER>0, min prior	-0.009	-0.003	-0.006	-0.012	-0.028	-0.023	-0.027	-0.032	-0.034	-0.034

Panel B: Equal-weighted Cryptocurrency

Prior	2013	2014	2015	2016	2017	2018	2019	2020	2021	End of Sample
Flat	1.021	-0.568	-0.413	-0.127	0.706	0.413	0.362	0.358	0.442	0.410
2	1.739	0.496	0.293	0.278	0.623	0.437	0.377	0.357	0.391	0.370
1	0.771	0.198	0.116	0.132	0.420	0.286	0.247	0.241	0.275	0.259
0	0.482	0.103	0.060	0.086	0.362	0.243	0.210	0.207	0.242	0.227
-1	0.173	0.001	-0.001	0.037	0.301	0.197	0.171	0.171	0.207	0.194
-2	-0.150	-0.108	-0.067	-0.015	0.238	0.149	0.130	0.134	0.170	0.158
-5	-1.255	-0.489	-0.297	-0.195	0.027	-0.013	-0.009	0.008	0.050	0.042
-10	-3.441	-1.270	-0.773	-0.562	-0.377	-0.328	-0.278	-0.234	-0.182	-0.182
-20	-9.244	-3.425	-2.097	-1.566	-1.403	-1.140	-0.973	-0.853	-0.763	-0.747
Max Gain	2.845	0.722	0.414	0.399	0.955	0.651	0.561	0.530	0.594	0.558
Optimal Prior	0.158	0.126	0.121	0.129	0.196	0.173	0.173	0.195	0.211	0.205
CER>0, max prior	0.328	0.259	0.247	0.279	0.455	0.408	0.410	0.429	0.477	0.467
CER>0, min prior	-0.015	-0.010	-0.009	-0.017	-0.053	-0.047	-0.048	-0.051	-0.061	-0.059

Table 8: Robustness and Extensions

This table considers a range of modifications to our baseline specification. Panel A considers different strengths of prior beliefs, ranging from observing three to 50 years of data before the beginning of the sample period. Panel B considers different baseline assets with which investors start, and compares cutoff beliefs for when it would not be worth adding cryptocurrency to the existing asset combination. Panel C varies investors' priors over the correlation of cryptocurrency with the market portfolio. Panel D explores how different prior beliefs about volatility affect the average weights on Bitcoin across different prior means, and Panel E explores how different prior beliefs about volatility affect the end-of-sample weights on Bitcoin across different prior means. Panel F reports the cutoff beliefs if investors ignore early data as being "unrepresentative" . Panel G reports the snapshot and accumulative cutoff prior for non-investment in Bitcoin under robust portfolio choice framework. Panel H summarizes weights on Bitcoin for different prior beliefs about the average monthly cryptocurrency excess return at the beginning of the sample period under robust portfolio choice framework. The sample consists of 106 monthly returns from May 2013 to February 2022.

Panel A: Strength of Prior Beliefs

Prior strength(Years)	3	5	10	30	50
BTC	-0.341	-0.206	-0.106	-0.038	-0.024
EWCrypto	-0.629	-0.382	-0.196	-0.072	-0.046

Panel B: Different Baseline Assets

	Market Only	3 Factors	4 Factors	4 Factors + MSCI
BTC	-0.106	-0.106	-0.106	-0.106
EWCrypto	-0.196	-0.199	-0.202	-0.199

Panel C: Different Priors about Correlations

Corr	0 (baseline)	0.1	0.2	0.3
BTC	-0.106	-0.091	-0.078	-0.065
EWCrypto	-0.196	-0.163	-0.137	-0.111

Panel D: Average Weight of Bitcoin for Different Volatility Priors

Volatility \BTC prior	2	1	0	-1	-2	-5	-10	-20
0.2x Sample	0.378	0.321	0.264	0.208	0.151	-0.017	-0.288	-0.792
0.5x Sample	0.168	0.146	0.124	0.103	0.081	0.016	-0.090	-0.294
Baseline	0.059	0.052	0.045	0.038	0.030	0.009	-0.026	-0.096
2x Sample	0.018	0.016	0.014	0.012	0.010	0.004	-0.006	-0.025
5x Sample	0.004	0.003	0.003	0.003	0.002	0.001	0.000	-0.004

Panel E: End-of-Sample Weight of Bitcoin for Different Volatility Priors

Volatility \BTC prior	2	1	0	-1	-2	-5	-10	-20
0.2x Sample	0.397	0.361	0.328	0.294	0.259	0.156	-0.008	-0.299
0.5x Sample	0.195	0.179	0.163	0.146	0.130	0.081	0.000	-0.150
Baseline	0.070	0.064	0.059	0.053	0.047	0.030	0.002	-0.055
2x Sample	0.021	0.019	0.018	0.016	0.015	0.010	0.002	-0.014
5x Sample	0.004	0.004	0.004	0.003	0.003	0.002	0.001	-0.002

Panel F: Prior for Cumulative Non-Investment at the end of Sample, with sequential dropping out data

Data Range	2014 -2022Feb	2015 -2022Feb	2016 -2022Feb	2017 -2022Feb	2018 -2022Feb	2019 -2022Feb	2020 -2022Feb	2021 -2022Feb
Btc-rf	-0.064	-0.070	-0.066	-0.058	-0.030	-0.037	-0.029	-0.015
ew-rf	-0.113	-0.118	-0.116	-0.107	-0.041	-0.049	-0.048	-0.035
vw-rf	-0.066	-0.071	-0.069	-0.061	-0.028	-0.038	-0.032	-0.018

Panel G: Snapshot and Cumulative Cutoff Priors for Bitcoin, with ambiguity aversion $\tau = 4$

	2013	2014	2015	2016	2017	2018	2019	2020	2021	End of Sample
Snapshot	-0.047	-0.042	-0.047	-0.055	-0.085	-0.077	-0.085	-0.098	-0.103	-0.104
Cumulative	-0.050	-0.050	-0.050	-0.055	-0.085	-0.085	-0.089	-0.098	-0.107	-0.107

Panel H: Optimal Bitcoin Weights with ambiguity aversion $\tau = 4$

Prior	Average	S.t.d	Lowest	Highest	Final	Fraction positive	Fraction above 0.5%	Fraction above 1%	Fraction above 2%	Fraction above 5%	First Date Weight is Positive	Mean of Leverage	S.t.d of Leverage
2	0.025	0.005	0.011	0.031	0.030	1.000	1.000	1.000	0.943	0.000	All above	0.421	0.016
1	0.022	0.005	0.006	0.028	0.027	1.000	1.000	0.943	0.642	0.000	All above	0.419	0.016
0	0.019	0.006	0.002	0.026	0.025	1.000	0.953	0.943	0.519	0.000	All above	0.415	0.017
-1	0.016	0.006	-0.002	0.023	0.023	0.953	0.943	0.943	0.321	0.000	2013-Oct	0.411	0.017
-2	0.013	0.006	-0.007	0.021	0.020	0.943	0.981	0.632	0.104	0.000	2013-Nov	0.407	0.018
-5	0.004	0.008	-0.020	0.014	0.013	0.642	0.585	0.245	0.000	0.000	2016-Jun	0.401	0.023
-10	-0.011	0.010	-0.041	0.001	0.001	0.104	0.632	0.453	0.113	0.000	2021-Feb	0.386	0.028
-20	-0.040	0.014	-0.081	-0.023	-0.023	0.000	1.000	1.000	1.000	0.274	All Below	0.353	0.031

7 Appendix - Towards a Theory of Bitcoin Prices

In this paper, we have been agnostic about different priors over cryptocurrency returns. We do not dispute the common basis for beginning with a discounted cash flows view of Bitcoin, which predicts a price of zero. Nonetheless, even pessimists ought to acknowledge that Bitcoin holders likely understand the lack of underlying cash flows, and their optimism does not stem from confusion about whether Bitcoin will produce a US-dollar-denominated dividend. Whatever is driving price changes is clearly something else. So what might that something else be?

We summarize some of the more credible cases for non-zero Bitcoin prices, noting that our arguments here are not original. These arguments are loose, polemic, incomplete, and proceed by analogy rather than as formal equilibrium models. Our results do not rely on their correctneess, however. They are starting points for further thinking, especially for readers who may be skeptical of our findings without some kind of economic basis. Something large is missing in our models of Bitcoin prices, and it seems appropriate to have some agnosticism over what that something is.

We consider four main (non-exhaustive) variants of arguments in favor of Bitcoin:

1. Bitcoin is a substitute for gold, and may ultimately replace it.

2. Bitcoin is a substitute for the US dollar, and may ultimately replace it.

3. Bitcoin is a substitute for both gold and the US dollar, and may ultimately replace both, because something gold-like will also replace the US dollar.

4. As a follow-on from any of the above, Bitcoin may become a substitute for either corporate liquidity holdings and/or central bank reserves.

We find #1 the most interesting and plausible, and potentially #4 as well. The "Bitcoin as gold" metaphor has a surprising amount to recommend. If one re-imagines a digitally stored and tradable version of gold for the 21st century, Bitcoin comes rather close. It is striking that many of the criticisms leveled at Bitcoin apply almost equally well to gold. Metal sits in the ground at various parts of the globe. Huge amounts of money, energy and resources are spent locating the metal, digging it out of the ground, purifying it, and then... putting it back into a different part of the ground, in the vaults of the Federal Reserve Bank of New York. [11]. Many of the bars have sat

[11] Gold equivalent to two years' worth of global annual production sit in the New York Fed vaults. Total central bank reserves plus bars and coins are equivalent to 27 years' worth of production in 2021. See

in the vault as long as any of us have been alive. We pose the challenge to economists - explain the role of those bars in the economy. If they had disappeared 30 years ago and nobody opened the vault, what would be different? They produce no cash flows, but can only be sold to a new buyer.

This suggests various questions. Is gold a bubble? It depends what is meant by this. There is some industrial and jewelry demand, so the equilibrium price will not be zero. It is less clear that these factors are the only or the main driver of gold prices. If the question means "is the price of gold substantially higher than it would be if all the central banks decided to stop holding gold?", then we think the answer is almost certainly "yes".[12]. Does this mean that the price of gold is about to collapse, or that gold is always a terrible investment? Not obviously, certainly not without a theory of why the price is high in the first place. Why gold, and not some other metal? Could the central banks all decide one day to switch to holding platinum, or silver, or molybdenum? Of course. Are they likely to? Not obviously. Why not? It is hard to say, but likely large factors include incumbency and the self-fulfilling liquidity that comes from many existing holders.

The largest difference between gold and Bitcoin is that gold has fundamental sources of demand from industrial uses and jewelry. A tomato also does not produce underlying cash flows, but it is something that consumers demand. Gold, unlike Bitcoin, has non-monetary users of the product. Industrial uses are the most obvious, but these cover a very small fraction of gold's history, so cannot be the explanation. Rather, the fundamental explanation that justifies the non-zero price is the wedding ring. Gold is valuable, so the theory goes, because people desire it for jewelry.

From this perspective, one underappreciated aspect of Nakamoto 2008 (on top of the blockchain, distributed trustless consensus and solving the double-spend problem) was realizing that the wedding ring may the least important aspect of gold. Indeed, economists may have causality backwards - people desire gold for wedding rings because it is valuable, not the other way around.

Rather, the most important aspect of gold may be its fixed supply, which means that price increases are not quickly reversed by expanded production. If some investors have extrapolative beliefs (such as diagnostic expectations in Bordalo, Gennaioli, and Shleifer 2018), these price increases may create more investor demand. There has to be some initial source of demand to get

`https://www.newyorkfed.org/aboutthefed/goldvault.html` for NY Fed numbers, `https://www.gold.org/goldhub/data/how-much-goldforgoldreserves`, and `https://www.statista.com/statistics/238414/global-gold-production-since-2005/` for production numbers

[12]The evidence for price pressure is considerable, even from predictable trades. See Hartzmark and Solomon 2021, Gabaix and Koijen 2021, and others.

the process going. But this could equally be true believers in Austrian economics and hard money, drug dealers, online poker players, or others wishing to avoid banks and financial reporting.

It seems useful that there is some narrative that makes the demand understandable - wedding rings are a convenient and reasonable explanation for gold, even if most trading comes from central banks and speculators. But at some point, most people buying Bitcoin are not drug dealers, or even thinking about drug dealers - they are just buying as a bet on Bitcoin prices.

Bitcoin is often cited as being a hedge asset, either against inflation or market downturns. These properties of returns have shown to not be fixed facts, as correlations of Bitcoin with market returns have increased since 2020. Rather, the hedge aspect is apt in a different sense. If you needs to flee the country you live in on 24 hours notice, to never return, with only what you can carry with you, there is no better asset for when you reach immigration at a new country. In answer to the question "are you transporting more than $10,000 of currency", it is unclear what that even means. The passphrase to a hardware wallet can be stored in one's head, and records of the coins live on computers all over the world. Bitcoin is only a hedge against intermediate disasters - personal catastrophe, or disaster in one country, where computers and the internet still operate. This is also true of gold - in a true post-apocalypse scenario, shiny metal is less useful than antibiotics, waterproof matches, water purification tablets, and guns.

The above argument elides over the biggest point - if fundamental value does not determine prices, what does? The somewhat tautological answer in Hartzmark and Solomon 2021 is "trades". Buyers and sellers submit limit orders. Intersections of these orders determine the price. While this is literally true, there is a sense in which the conclusions from the idea are quite surprising. One can imagine various hypotheses for prices

H0: Prices are equal to fundamental value - the asset pricing market efficiency view

H1: Prices are not always equal to fundamental value, and may deviate when psychological biases or frictions lead to errors - the standard behavioral finance view

H2: Prices can be anything, unless some specific force constrains them to a particular range.

H2 has a surprising amount to recommend about it. Fundamental value is one force constraining prices, inasmuch as when prices deviate too far from fundamental value, there are profitable trading strategies that do not depend on other traders changing their minds. When prices are below fundamental value, it is profitable to buy the asset, hold it, and collect the cash flows. If one is patient and not subject to capital withdrawals, this strategy has considerable appeal, but may take a long time to work. If the asset trades far above fundamental value, it is profitable to buy the underlying assets of the project, create a new version of it, and sell equity in it, though this has large implementation risks. For Bitcoin, one can fork the code, and even the state of the ledger (as Bitcoin Cash did), but one faces an uphill battle to overcome incumbency. In supporting a forked version one is spending valuable hash power on a coin that will likely fail. Most other strategies using fundamental value require other traders updating their beliefs towards fundamental value. Without a theory of what the mistake is, it is hard to know when or if it will be corrected. These trades are both a bet on the existence of mispricing, and a bet on when it will correct itself.

If prices change because of trades and price pressure, what are the important aspects of Bitcoin? The fixed supply is obviously important, as is the general difficulty of shorting Bitcoin (now somewhat eased). Another is the strong culture by the largest and earliest wallet owners to hold and never sell ('hodl', in the neologism of Bitcoiners). Incidentally, this aspect rather resembles central banks attitudes to gold. Bitcoin's much greater uncertainty, and much more fluctuating investor base, make it much more volatile. Because there simply are no cash flows, all trading becomes a coordination game. This analysis predicts that the volatility of Bitcoin need not be something that will settle down in the near term, even if it is successful. In other words, the stability of gold is a function of who holds it and why, not a fixed property of the metal.

We briefly turn to arguments #2 and #3. Both involve a belief in Austrian economics, that governments will realize the virtues of fixed monetary policy, or will be forced into it. Such arguments about monetary theory are beyond the scope of this paper, but we are somewhat skeptical. The scaling issues of the ledger size and transaction rate limit are well known, although Layer 2 solutions (such as the lightning network for Bitcoin, and a variety of platforms for Ethereum) may circumvent this. More importantly, it seems likely that if Bitcoin were credibly likely to replace US dollars, it would be viewed as a major threat to the US government, as it threatens the ability to raise debt through printing money. Despite skepticism, we remain agnostic - if Bitcoin has taught us anything, it is that economists' models of money seem to be missing something large.

Finally, #4 has become especially pertinent following the US and EU decisions to seize the assets of the Russian Central Bank during the war in Ukraine in 2022. This event highlighted that it is not only Bitcoin that exists primarily as numbers stored on a computer, but most financial assets like stocks and currencies. The only difference is that for most assets, the computers are amenable to the control of foreign governments. If the question is "what is the largest, most liquid electronically traded asset that cannot be seized by foreign governments during a crisis?", the answer is "Bitcoin". It seems plausible that more central banks may eventually purchase cryptocurrency just as hard asset reserves, rather than for everyday currency purposes like El Salvador .

These arguments, if informal, suggest some out of sample predictions. Firstly, incumbency is a large advantage in the coordination game of which asset is to be the gold-like hedge asset. Bitcoin may be supplanted by something like Ethereum which offers technological improvements and features that Bitcoin lacks (as Bitcoin improved on gold). But the logic would predict that Bitcoin is unlikely to be replaced by other smaller and less liquid "pure exchange" objects that are substantially Bitcoin substitutes, such as Dogecoin or others. An important metric for success in this respect is total market capitalization. Bitcoin has been the largest cryptocurrency at all points. If it loses this status to another coin that offers better technology, this would be a negative sign, as it is not clear what force would cause the coordination focus to move back the other way.

Second, if one believes that Bitcoin is a potential gold substitute, the gold market capitalization of $11.4 trillion is much greater than the Bitcoin market capitalization ($387 billion as of August 29th, 2022). There is no guarantee that all the gold demand will move to Bitcoin, especially demand from central banks and institutions. Nonetheless, it provides some basis for possible belief in scenarios where the Bitcoin price still rises considerably relative to its current level.

Third, the culture of "hodling" among old, large holders of Bitcoin is likely important. It is not clear when or why these holders might change their mind and sell (and such transactions would likely reveal their identity). But if such selling occurs, it would be considerable negative news, given the large amounts of Bitcoin they hold, and the potential for self-fulfilling beliefs to unravel.

Fourth, new technological developments are interesting in terms of the likely distribution of buy and sell orders they generate. For instance, a Bitcoin ETF is unclear in its likely effect. On the one hand, it makes it easier for investors with self-directed IRAs to purchase Bitcoin, leading to predictions of price increases. On the other hand, it is also easier to short a Bitcoin ETF than to take a short futures position. As a consequence, the net impact on prices is not clear.

Finally, there are many ways the self-reinforcing cycle of belief in Bitcoin's value as an asset could unwind. All of the above is not an argument that Bitcoin *will* succeed, merely that it *might*.

Internet Appendix for:

"The Cryptocurrency Participation Puzzle"

This appendix presents our methods and estimates from robustness tests and extensions. Section 1 discusses the Bayesian methods used to calculate the posterior distribution. Section 2 discusses robust portfolio choice methods used to consider model misspecification. Section 3 provides estimates from robustness tests and extensions.

A Bayesian Methods

In addition to the basic portfolio choice setup described previously, we now describe the mechanics of calculating the posterior distributions of the key parameters. Assume the (excess) returns of the N risky asset follow the following data generating process (DGP):

$$r = \mu + \epsilon,\ \epsilon \sim N(0, V) \tag{A1}$$

In our study, there are $N = 6$ assets, namely: the cryptocurrency asset-rf, mkt-rf, SmB, HmL, UmD and MSCIexUS-rf.

The informative priors on Mean μ and Covariance V are given by:

$$\mu \sim N(\mu_0, \Omega V \Omega),\ \Omega = \begin{bmatrix} \rho_1 & \dots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \dots & \rho_N \end{bmatrix} \tag{A2}$$

$$V \sim IW(V_0, h) \tag{A3}$$

Following the standard empirical Bayesian approach, the prior on the mean μ, denoted as μ_0, is set to the sample mean in the pre-period preceding our portfolio choice. Intuitively, when this initial sample period is shorter, the investor would assign a weaker prior since there is less data to pin down the prior. Hence, we set ρ^2 as the inverse of the number of months in the initial sample period. This way, the shorter the initial sample data is, the larger the value of ρ is, and consequently the weaker the prior is. As $1/\rho^2$ corresponds to the strength of the prior mean, we set $1/\rho^2$ for Mkt-rf, SMB, HML, UMD as 1042, since we have 1042 months of pre-sample data for these assets. We set $1/\rho^2 = 520$ for MSCIexUS, as we have 520 months of data for it. Since we do not have pre-sample data for cryptocurrency assets, we select different values for $1/\rho^2$ to reflect different prior strengths. In the baseline case, we set $1/\rho^2 = 120$, which corresponds to an initial cryptocurrency sample data of 10 Years. For the flat prior, we set $1/\rho^2 = 1e-6$.

Moreover, when a risky asset i is more volatile, there is more parameter uncertainty about its mean μ_i. Hence, when the variance σ_i of the returns on a risky asset i is higher, we assume a weaker prior on μ_{i0} to reflect the higher parameter uncertainty.

Similarly, following the standard empirical Bayesian approach, V_0 is based on the sample covariance matrix using pre-period sample data. We set $V_0 = (h-k)V_{prior}$, where $V_{prior} \in R^{k\times k}$ and is represented by

$$V_{prior} = \begin{bmatrix} \sigma^2_{cry} & 0 & 0 & 0 & 0 \\ 0 & & & & \\ 0 & & V_{factors} & & \\ 0 & & & & \\ 0 & & & & \end{bmatrix} \tag{A4}$$

where $V_{factors}$ is the sample covariance matrix of the five factors, namely, mkt =Mkt-rf, SMB, HML, UMD and MSCIexUS-rf based on 520 months of data prior to our portfolio choice analysis. Then, for these five factors, the expected covariance matrix is the same as the pre-period sample covariance matrix. In the baseline case, we set σ^2_{crypto} as $150, 170, 625$ times σ^2_{mkt} for the returns on Bitcoin (BTC), the value-weighted crypto market portfolio (VW100) and the equal-weighted crypto market portfolio (EW100), respectively. These ratios are based on the empirical sample ratios between the variance of the cryptocurrency assets and the stock market variance σ^2_{mkt}. We consider alternative priors on the variances of cryptocurrency assets in robustness tests.

Furthermore, in the baseline case we set the prior correlation between cryptocurrencies and the market portfolio as zero. We do so because the correlation was roughly zero in the decade following the introduction of Bitcoin in 2009. For robustness, we re-estimate the analyses using alternative correlation priors.

$$V_{prior} = \begin{bmatrix} \sigma^2_{cry} & \rho\sigma_{cry}\sigma_{mkt} & 0 & 0 & 0 \\ \rho\sigma_{cry}\sigma_{mkt} & & & & \\ 0 & & V_{factors} & & \\ 0 & & & & \\ 0 & & & & \end{bmatrix} \tag{A5}$$

Using the above priors, we derive the posterior distribution. First, we derive the following density functions for the priors:

$$f_0(\mu|V) \propto |\Omega V \Omega|^{-\frac{1}{2}} \exp\Big\{ -\frac{1}{2}\left[\mu - \mu_0\right]' \frac{1}{\Omega} V^{-1} \left[\mu - \mu_0\right] \Big\}$$
$$p_0(V) \propto |V|^{-\frac{h+n+1}{2}} \exp\Big\{ -\frac{1}{2} tr\left[V_0 V^{-1}\right] \Big\}$$

Second, we derive the posterior joint distribution of μ, V proportional to $f(R|\mu, V)f_0(\mu|V)f_0(V)$,

$$
\begin{aligned}
f(\mu, V|r_1, \ldots, r_T) \propto & |\Omega V \Omega|^{-\frac{1}{2}} \exp\Big\{ -\frac{1}{2}[\mu - \mu_0]' \frac{1}{\Omega} V^{-1} \frac{1}{\Omega} [\mu - \mu_0] \Big\} \\
& \times |V|^{-\frac{h+n+1}{2}} \exp\Big\{ -\frac{1}{2} tr\left[V_0 V^{-1}\right] \Big\} \\
& \times |V|^{-\frac{T}{2}} \exp\Big\{ -\frac{1}{2} \left[\sum_{t=1}^{T} (r_t - \mu)' V^{-1} (r_t - \mu)\right] \Big\}
\end{aligned}
$$

where $\frac{1}{\Omega}$ stands for Ω^{-1}.

To implement Gibbs Sampling, we derive the posterior conditional distribution $f(\mu|V, R)$ and $f(V|\mu, R)$ as follows

$$
\begin{aligned}
f(\mu|(r_1, \ldots, r_T), V) \propto & \exp\Big\{ -\frac{1}{2}[\mu - \mu_0]' \frac{1}{\Omega} V^{-1} \frac{1}{\Omega} [\mu - \mu_0] - \frac{1}{2} \sum_{t=1}^{T} (\mu - r_t)' V^{-1} (\mu - r_t) \Big\} \\
\propto & \exp\Big\{ -\frac{1}{2} \left[\mu' \left(\frac{1}{\Omega} V^{-1} \frac{1}{\Omega} + T V^{-1} \right) \mu - 2 \left(\mu_0' \frac{1}{\Omega} V^{-1} \frac{1}{\Omega} + (\sum_t r_t)' V^{-1} \right)' \mu \right] \Big\}
\end{aligned}
$$

Therefore, conditional on V, μ is normally distributed, with the following mean:

$$
\left(\frac{1}{\Omega} V^{-1} \frac{1}{\Omega} + T V^{-1} \right)^{-1} \left(\frac{1}{\Omega} V^{-1} \frac{1}{\Omega} \mu_0 + V^{-1} (\sum_t r_t) \right) = \left(\frac{1}{\Omega} V^{-1} \frac{1}{\Omega} + T V^{-1} \right)^{-1} \left(\frac{1}{\Omega} V^{-1} \frac{1}{\Omega} \mu_0 + T V^{-1} \bar{r}_T \right)
$$

When $\Omega = \rho I$, this equation collapses into:

$$
\frac{T}{T + 1/\rho^2} \bar{r}_T + \frac{1/\rho^2}{T + 1/\rho^2} \mu_0 \tag{A6}
$$

Note that $1/\rho^2$ intuitively measures how much investors trust the prior μ_0, assuming they observe data with $1/\rho^2$ periods and mean μ_0. As $\rho \to 0$, investors will trust μ_0 as if it has been observed over an infinitely long period, and consequently the sample average becomes the estimator for the population mean. In addition, μ has a conditional covariance matrix $\left(\frac{1}{\Omega} V^{-1} \frac{1}{\Omega} + T V^{-1}\right)^{-1}$.

Moreover, for the posterior distribution of V, we have that $V|\mu, R$ is IW distributed, due to the following reasoning:

$$
\begin{aligned}
&f(V|(r_1,\ldots,r_T),\mu) \propto |\Omega V\Omega|^{-\frac{1}{2}}|V|^{-\frac{h+N+1}{2}}\exp\Big\{-\frac{1}{2}tr\left[V_0^{-1}V\right]\Big\} \\
&\times |V|^{-\frac{T}{2}}\exp\left[-\frac{1}{2}\sum_{t=1}^{T}(r_t-\mu)'V^{-1}(r_t-\mu)-\frac{1}{2}\left(\mu-\mu_0\right)'\frac{1}{\Omega}V^{-1}\frac{1}{\Omega}\left(\mu-\mu_0\right)\right] \\
&\propto |V|^{-\frac{2+N+T+h}{2}}\exp\Big\{-\frac{1}{2}tr\left[V_0V^{-1}\right]\Big\}\exp\Big\{-\frac{1}{2}tr\left[\left(\sum_t(r_t-\mu)(r_t-\mu)'\right)V^{-1}\right]\Big\} \\
&\times \exp\Big\{-\frac{1}{2}tr\left[\frac{1}{\Omega}\left(\mu-\mu_0\right)\left(\mu-\mu_0\right)'\frac{1}{\Omega}V^{-1}\right]\Big\} \\
&\propto |\rho V|^{-\frac{2+N+T+h}{2}}\exp\Big\{-\frac{1}{2}tr\left[\left(V_0+\sum_t(r_t-\mu)(r_t-\mu)'+\frac{1}{\Omega}\left(\mu-\mu_0\right)\left(\mu-\mu_0\right)'\frac{1}{\Omega}\right)V^{-1}\right]\Big\}
\end{aligned}
$$

Hence, $V|\mu,R \sim IW(V_0+\sum_t(r_t-\mu)(r_t-\mu)'+\Omega^{-1}\left(\mu-\mu_0\right)\left(\mu-\mu_0\right)'\Omega^{-1}, T+h+1)$.

Given the conditional posterior distributions, $\mu|V,R$ and $V|\mu,R$, we implement the Gibbs sampling process to obtain the estimated values of the predictive means and covariances of the six assets employed in our study. To be more precise, after each sampling of μ, V, we sample predictive return from normal distribution with mean μ and covariance matrix V. Therefore, we can estimate the predictive density of R_{t+1}.

B Robust Portfolio Choice

Following Hansen and Sargent (2001), Anderson, Hansen and Sargent (2003), and Anderson and Chen (2016), we allow agents to worry about model misspecification by considering perturbations to the probability density of asset returns that can decrease utility. Equivalently, people may cast doubt about a single probability law to describe the distribution of the relevant random variables. Although there are multiple ways to perturb distributions, we may assume that there is an adversarial agent acting against us. Given our portfolio choice, the adversarial agent perturbs the distribution to decrease our utility but incurs some perturbation cost. Meanwhile agents'fears about model specification doubts can be measured by ambiguity aversion. Under the adversarial agent setting, ambiguity aversion can be viewed as equivalently to adversarial agent's perturbation cost. If the adversarial agent has smaller cost to perturb distributions, they can perturb the distribution in a larger distribution space. Hence, the higher perturbation possibility investor should worry about or have higher level of ambiguity aversion. Taking the existence of this adversarial agent into consideration, investors choose their portfolio to maximize utility that is minimized by the adversary agent, or we are trying to solve a max-min or min-max question. This robust portfolio choice problem is similar spiritually to recent popular Generative Adversarial Networks in CS literature(Goodfellow et.al.(2014)).

According to investors' belief, we have a density function for future return R, let $z = R - E[R]$.

Denote the p.d.f of z as $f(z)$, and perturbed distribution of z as $f(z)\rho(z)$. The robust portfolio choice problem can be set up as following:

$$\max_{\phi}\min_{\rho}\left[\int_{z}\left[\phi'\left(\mu+z\right)-\frac{\gamma}{2}\left(\phi' z\right)^{2}\right]f\left(z\right)\rho(z)dz \right. \\ \left. +\frac{1}{\tau}\int\rho\left(z\right)f\left(z\right)\log\rho\left(z\right)dz\right] \tag{A7}$$

s.t.

$$\int\rho\left(z\right)f\left(z\right)dz=1 \tag{A8}$$

where the second term of equation A7 stands for the relative entropy(Kullback-Leibler divergence) between $\rho(z)f(z)$ and $f(z)$, measuring the discrepancy between the perturbed distribution and original one. And τ measures the ambiguity aversion level of investor, i.e. $1/\tau$ measures the cost of perturbing distribution.

Variational method and constraint A8 delivers the optimal condition for $\rho(z)$:

$$\varrho^{*}\left(z\right)=\frac{\exp\left[-\tau\phi' z+\frac{\theta\tau}{2}\left(\phi' z\right)^{2}\right]}{E\left(\exp\left[-\tau\phi' z+\frac{\theta\tau}{2}\left(\phi' z\right)^{2}\right]\right)} \tag{A9}$$

Substituting equation A9 into equation A7, we can write the optimal portfolio choice problem as following

$$\max_{\phi}\left(\phi'\mu^{*}-\frac{1}{\tau}\log\int\exp\left[-\tau\phi' z+\frac{\theta\tau}{2}\left(\phi' z\right)^{2}\right]f\left(z\right)dz\right) \tag{A10}$$

Though we have no analytical form for $f(z)$, which corresponds to predictive density of future returns, we can draw random sample from $f(z)$ and solve out the optimal portfolio choice according to equation A10.

Previous literature shows ambiguity aversion acts as an extra part of risk aversion(Trojani and Vanini (2004)). To offer an intuitive illustration for effects of ambiguity aversion and our following empirical results, we use Taylor expansions of equation A10. When our portfolio ϕ is close to zero, we have

$$
\begin{aligned}
&\max_{\phi}\left(\phi'\mu^* - \frac{1}{\tau}\log\int \exp\left[-\tau\phi' z + \frac{\theta\tau}{2}\left(\phi' z\right)^2\right] f\left(z\right)dz\right)\\
&\approxeq \max_{\phi}\left(\phi'\mu^* - \frac{1}{\tau}\log\int (1-\tau\phi' z + \frac{\tau(\theta+\tau)(\phi' z)^2}{2})f(z)dz\right)\\
&= \max_{\phi}\left(\phi'\mu^* - \frac{1}{\tau}\log(1+\frac{\tau(\theta+\tau)\phi'\Sigma\phi}{2})\right)\\
&\approxeq \max_{\phi}\left(\phi'\mu^* - \frac{1}{\tau}(\frac{\tau(\theta+\tau)\phi'\Sigma\phi}{2})\right)\\
&= \max_{\phi}\left(\phi'\mu^* - (\frac{(\theta+\tau)\phi'\Sigma\phi}{2})\right)
\end{aligned}
$$

where the second line follows Taylor expansion, omitting $o((\phi' z)^2)$. The third line follows the definition of covariance matrix of z. The fourth line follows $\lim_{x\to 0}\frac{\log(1+x)}{x} = 1$. As one can tell from the final result, ambiguity aversion coefficient act as an extra term of risk aversion coefficient.

C Robustness and Extensions

In this section, we present a number of additional analyses.

Figure A1 plots the continuous ex-ante certainty equivalent of returns gains from access to cryptocurrency, shown over time for various different priors. Panel A plots these for Bitcoin, while Panel B plots these for the equal-weighted cryptocurrency portfolio. These numbers are computed assuming that short sales are possible for Bitcoin after December 2017, but not possible for the equal-weighted portfolio. This is analogous to the results in Table 4 of the paper.

Figure A2 is similar, except that it plots the ex-post CER gains. It is the continuous version of Table 7 of the paper.

Figure A3 plots how the end-of-sample cutoff priors for non-investment vary according to the strength of priors. The latter is plotted in terms of the number of years of data the investor is presumed to have seen before the sample begins. Panel A shows this for Bitcoin, Panel B shows this for an equal-weighted cryptocurrency portfolio. These are the continuous version of Table 8 Panel A.

Figure A4 shows how end-of-sample cutoff mean priors for non-investment vary with the investor's priors about correlations between cryptocurrency and the equity market. Priors over correlations are considered between values of -0.5 and 0.5. Panel A shows the mean cutoff beliefs for Bitcoin, Panel B shows this for an equal-weighted cryptocurrency portfolio. These are the continuous versions of Table 8 Panel C.

Table A1 shows the priors required for non-investment in cryptocurrency at various points in the sample that represent local peaks and troughs of cryptocurrency prices. These are shown for Bitcoin (rows 1 and 2), equal-weighted cryptocurrency (rows 3 and 4), and value-weighted cryptocurrency (rows 5 and 6). Rows 1, 3 and 5 are priors for non-investment at that particular point in time, while rows 2, 4 and 6 are for zero or negative desired weights up to that point in time.

Table A2 presents the main results of the paper for the value-weighted cryptocurrency portfolio, analogous to those in the main text for the equal-weighted portfolio and bitcoin. Panel A presents summary statistics for the optimal portfolio weights at different points in time (the equivalent of Table 3). Panel B shows the ex-ante certainty equivalent gains at different points in time (the equivalent of Table 4). Panel C examines various levels of investment costs, and shows which priors are deterred from ever investing up to that point (the equivalent of Table 6 Panel A). Panel D shows a range a priors, and computes which costs would deter investment up to that point (the equivalent of Table 6 Panel C). Panel E shows the ex-post certainty equivalent gains from access to cryptocurrency at each point (the equivalent of Table 7).

Table A3 shows the ex-ante gains in Sharpe Ratios from having access to cryptocurrency, for various different priors, for bitcoin (Panel A), equal-weighted cryptocurrency (Panel B), and value-weighted cryptocurrency (Panel C).

Table A4 shows how certainty equivalent gains vary with transaction costs. Panel A examines ex-ante CER gains, while Panel B examines ex-post CER gains. These are computed at the end of the sample for annual costs of 2%, 5%, 10% and 20%, for each cryptocurrency portfolio, and for the range of priors examined in other tables.

Table A5 presents the cutoff priors for non-investment under ambiguity aversion of $\tau = 4$, computed for the equal-weighted and value-weighted cryptocurrency portfolios. This is shown for non-investment at each time (Panel A), and no positive weights up to that point (Panel B). This is the equivalent of Table 8 Panel G.

Table A6 presents summary statistics for desired portfolio weights under ambiguity aversion of $\tau = 4$. Panel A shows equal-weighted cryptocurrency, and Panel B shows value-weighted cryptocurrency. This is the equivalent of Table 8 Panel H.

Table A7 shows how desired weights vary according to perceived volatility levels. Panels A and B examine the equal-weighted cryptocurrency portfolio, with Panel A showing average weights and Panel B showing end-of-sample weights. Panels C and D show the same results (average and end-of-sample respectively) for the value-weighted cryptocurrency portfolio. These results are the equivalent of Table 8 Panels D and E.

Figure A1: The Time Series of Ex-Ante Certainty Equivalent Gains from Cryptocurrency

This figure plots the time series of certainty equivalent of return (CER) gains from adding cryptocurrencies to investors' existing portfolios for different prior beliefs about the average monthly cryptocurrency return at the beginning of the sample period. The reported values equal the difference between the CER of the baseline market portfolio that excludes cryptocurrency and the CER of the optimal portfolio that combines the market portfolio and cryptocurrency, assuming that investors have a constant relative risk aversion of 3. We assume that investors can short Bitcoin starting December 2017, when the Chicago Mercantile Exchange (CME) introduced future contracts on Bitcoin, but cannot short the equal-weighted cryptocurrency portfolio throughout the sample period. Panel A shows the CER gains for Bitcoin, whereas Panel B shows the CER gains for the equal-weighted cryptocurrency portfolio. The sample consists of 106 monthly returns from May 2013 to February 2022.

Panel A: Bitcoin

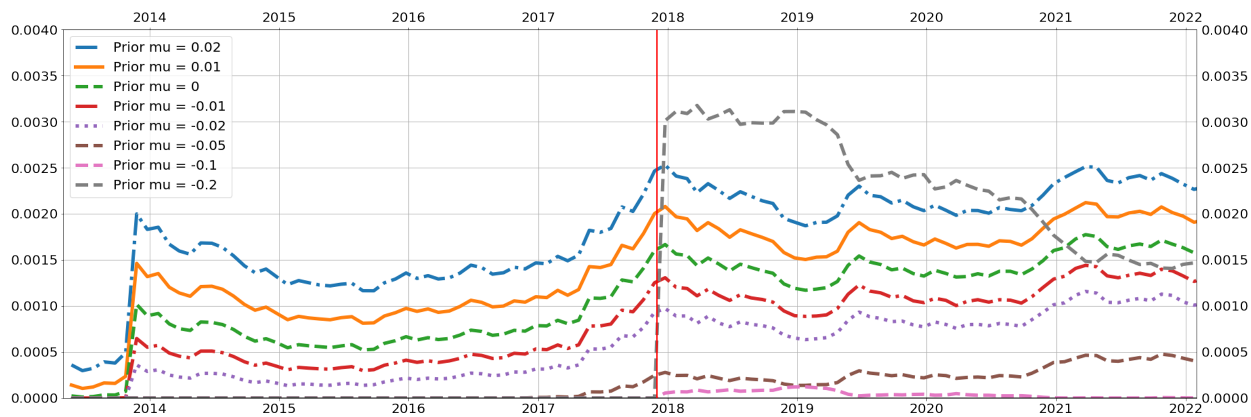

Panel B: Equal-weighted Cryptocurrency Portfolio

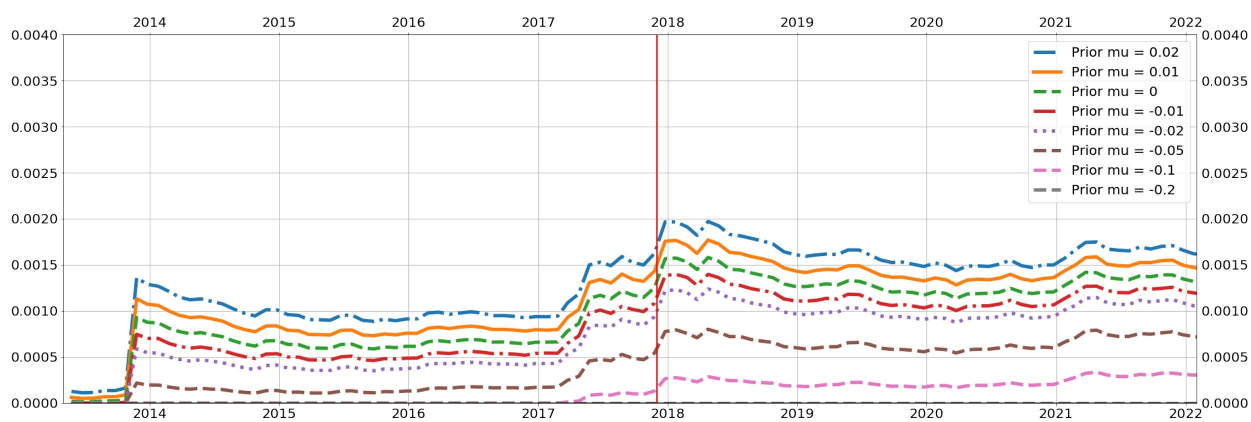

Figure A2: The Time Series of Ex-Post Certainty Equivalent Gains from Cryptocurrency

This figure plots the time series of ex-post certainty equivalent return (CER) gains from adding cryptocurrencies to investors' existing portfolios for different prior beliefs about the average monthly cryptocurrency return at the beginning of the sample period. Investors assess ex-post performance on a distributional basis, assuming that the distribution of realized returns up to that point (from whatever series of weights was chosen) were to continue indefinitely. The reported values equal the difference between the ex-post CER of the baseline market portfolio that excludes cryptocurrency and the ex-post CER of the optimal portfolio that combines the market portfolio and cryptocurrency. Investors are assumed to have a constant relative risk aversion of 3. We assume that investors can short Bitcoin starting December 2017, when the Chicago Mercantile Exchange (CME) introduced future contracts on Bitcoin, but cannot short the equal-weighted cryptocurrency portfolio throughout the sample period. Panel A shows the ex-post CER gains for Bitcoin, whereas Panel B shows the ex-post CER gains for the equal-weighted cryptocurrency portfolio. The sample consists of 106 monthly returns from May 2013 to February 2022.

Panel A: Bitcoin

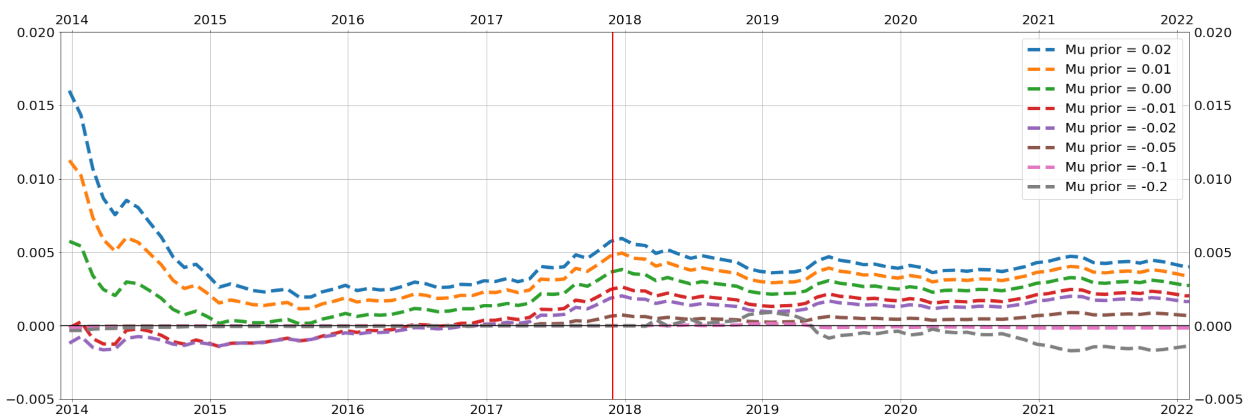

Panel B: Equal-weighted Cryptocurrency Portfolio

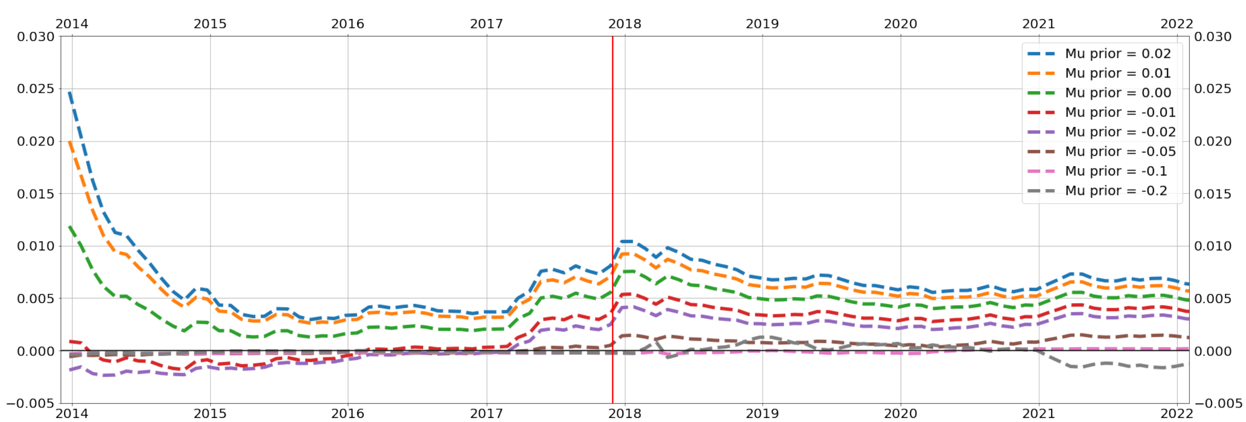

Figure A3: The Strength of Prior Beliefs and Non-Investment in Cryptocurrency

This figure plots the cutoff prior beliefs about the average monthly cryptocurrency return at the beginning of the sample period required for non-investment in cryptocurrency by the end of the sample period, as a function of the strength of the prior (ranging from 5 to 100 years of pre-sample data observed by investors). Panel A shows the cutoff beliefs for Bitcoin, whereas Panel B shows the cutoff beliefs for the equal-weighted cryptocurrency portfolio. The sample consists of 106 monthly returns from May 2013 to February 2022.

Panel A: Bitcoin

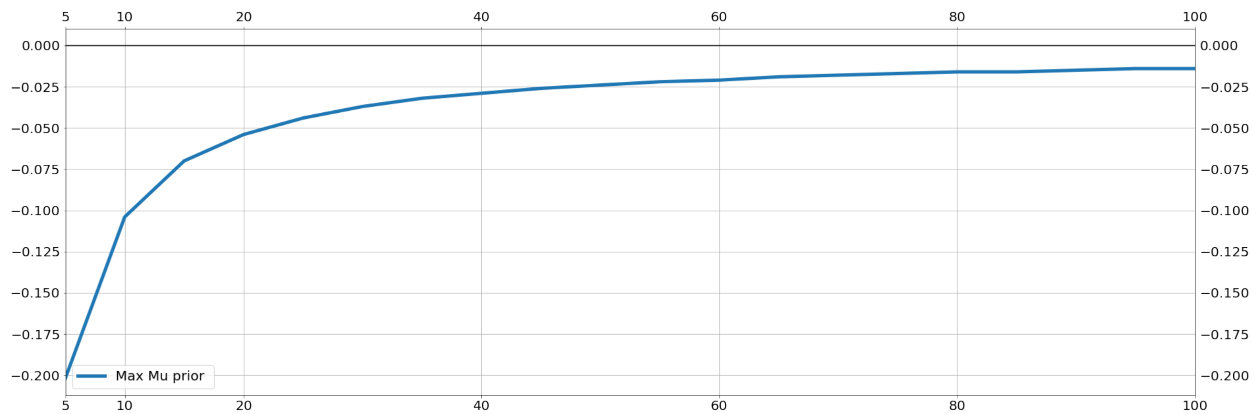

Panel B: Equal-weighted Cryptocurrency Portfolio

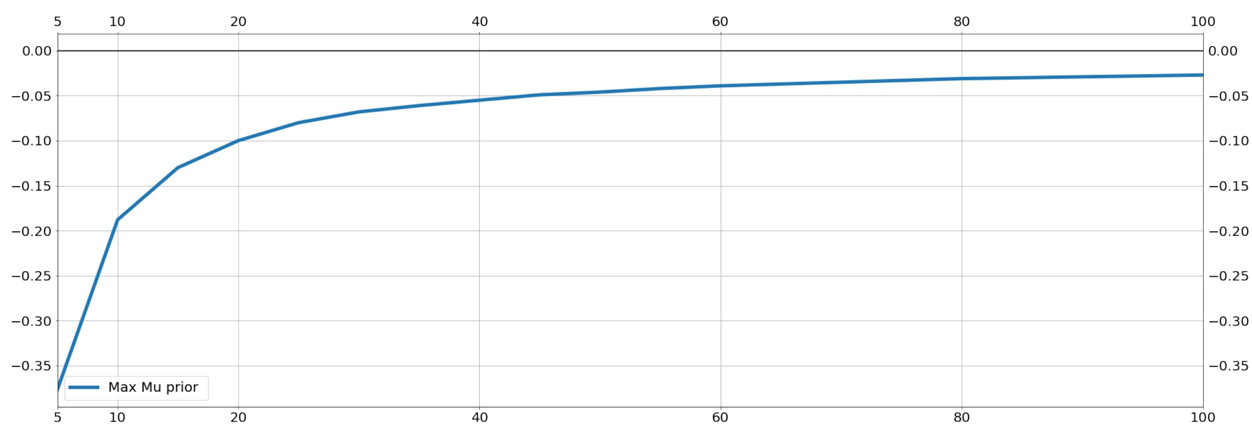

Figure A4: Prior Beliefs about Correlations and Non-Investment in Cryptocurrency

This figure plots the cutoff prior beliefs about the average monthly cryptocurrency return at the beginning of the sample period required for non-investment in cryptocurrency by the end of the sample period, as a function of prior beliefs about the correlation between cryptocurrencies and the market portfolio. Panel A shows the cutoff beliefs for Bitcoin, whereas Panel B shows the cutoff beliefs for the equal-weighted cryptocurrency portfolio. The sample consists of 106 monthly returns from May 2013 to February 2022.

Panel A: Bitcoin

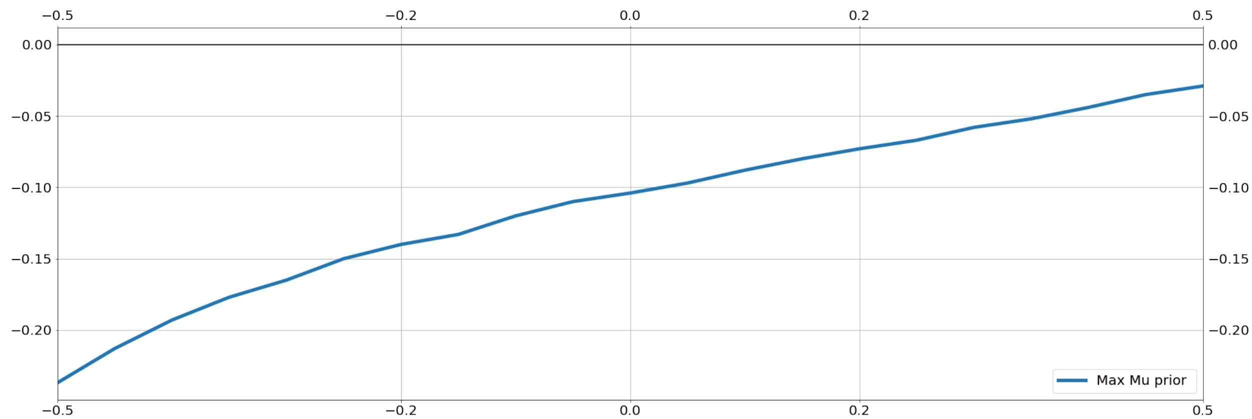

Panel B: Equal-weighted Cryptocurrency Portfolio

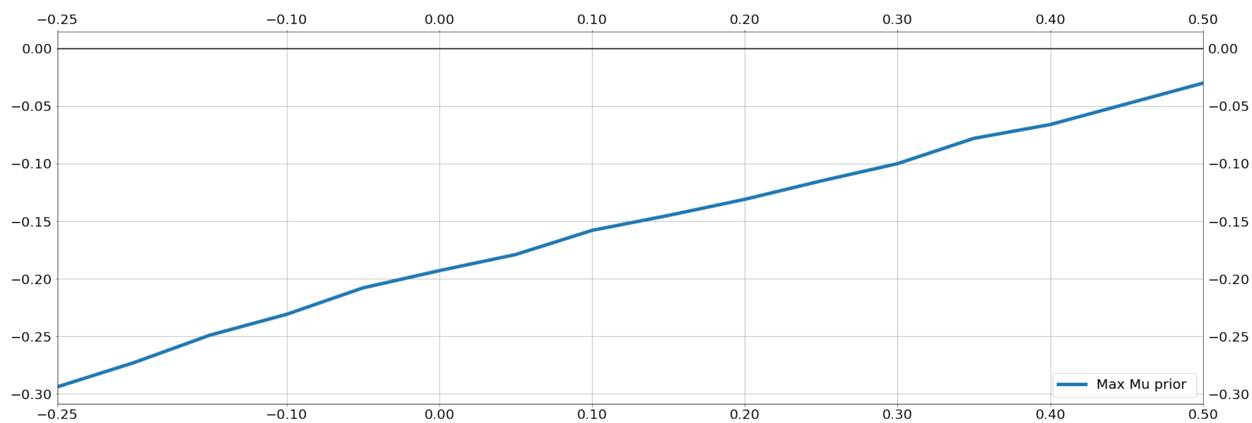

Table A1: Peaks and Troughs

This table reports cutoff prior beliefs about the average monthly cryptocurrency return at the beginning of the sample period that would make an investor not invest in cryptocurrencies. Each row corresponds to a different trough or peak months during the sample period, and the columns correspond to different cryptocurrency portfolios (Bitcoin, Equal-weighted cryptocurrency portfolio, Value-weighted cryptocurrency portfolio). The odd columns calculate the cutoff prior belief that leads to no cryptocurrency investment on a specific date, whereas the even columns calculate the cutoff belief that leads to no cryptocurrency investment at any point prior to the date. If the priors are above (below) the cut-ff level, then investors should long (short) on a specific date (odd columns) or at some point prior to the date (even columns). The calculations assume the following: (1) Investors start with the CRSP value-weighted market portfolio as a base asset and consider adding cryptocurrencies to their portfolios; (2) Investors observed ten years of data with a mean equal to their prior mean before the beginning of the sample period; (3) The variance of cryptocurrency returns approximately equals their ex-post variance – 150 times the market variance for bitcoin, 170 times the market variance for the value-weighted cryptocurrency portfolio, and 625 times the market variance for the equal-weighted portfolio; (4) Investors believe cryptocurrency to be uncorrelated with the market portfolio. The sample consists of 106 monthly returns from May 2013 to February 2022.

	BTC		EW		VW	
	Snap-shot	Cumu-lative	Snap-shot	Cumu-lative	Snap-shot	Cumu-lative
Jun-13 (2013 Trough)	-0.005	-0.007	-0.011	-0.013	-0.005	-0.007
Sep-13 (2013 Local Trough)	-0.009	-0.009	-0.015	-0.015	-0.009	-0.009
Nov-13 (2013-2014 Peak)	-0.050	-0.050	-0.097	-0.097	-0.052	-0.052
Apr-14 (Mt Gox Trough)	-0.044	-0.050	-0.091	-0.097	-0.046	-0.052
Jun-14 (2014 Peak)	-0.047	-0.050	-0.092	-0.097	-0.049	-0.052
Jan-15 (2014-2015 Trough)	-0.040	-0.050	-0.088	-0.097	-0.042	-0.052
Aug-15 (2015 Local Trough)	-0.041	-0.050	-0.090	-0.097	-0.043	-0.052
Jan-16 (2016 Trough)	-0.046	-0.050	-0.093	-0.097	-0.047	-0.052
Jun-16 (2016 Local Trough)	-0.051	-0.051	-0.102	-0.102	-0.053	-0.053
Aug-16 (2016 Local Trough)	-0.050	-0.051	-0.100	-0.102	-0.051	-0.053
Mar-17 (2017 Local Trough)	-0.056	-0.058	-0.113	-0.113	-0.059	-0.059
Dec-17 (Pre-2017 Peak)	-0.085	-0.085	-0.171	-0.171	-0.091	-0.091
Jan-19 (19-20 Trough)	-0.076	-0.085	-0.161	-0.175	-0.080	-0.091
Jun-19 (2019 Peak)	-0.090	-0.090	-0.169	-0.175	-0.091	-0.091
Mar-20 (2020 Trough)	-0.084	-0.090	-0.163	-0.175	-0.087	-0.091
Sep-20 (2020 Local Trough)	-0.088	-0.090	-0.171	-0.175	-0.091	-0.092
Mar-21 (2021 Local Peak)	-0.103	-0.103	-0.192	-0.192	-0.106	-0.106
May-21 (2021 Local Trough)	-0.100	-0.103	-0.191	-0.193	-0.105	-0.107
Aug-21 (2021 Local Peak)	-0.102	-0.103	-0.194	-0.194	-0.108	-0.108
Sep-21 (2021 Local Trough)	-0.102	-0.103	-0.193	-0.194	-0.107	-0.108
Oct-21 Peak	-0.106	-0.106	-0.195	-0.195	-0.110	-0.110
Jan-22 (2022 Local Trough)	-0.102	-0.106	-0.192	-0.196	-0.106	-0.110
Feb-22 (End of Sample)	-0.103	-0.106	-0.192	-0.196	-0.107	-0.110

Table A2: The Value-Weighted Cryptocurrency portfolio

This table re-estimates the main analyses for the value-weighted cryptocurrency portfolio. Panel A reports cryptocurrency portfolio weights for different prior beliefs about the average monthly cryptocurrency excess return at the beginning of the sample period. Panel B reports the monthly certainty equivalent return (CER) gains, in percentage points, from adding cryptocurrency to investors' existing portfolios. Panel C reports the range of prior means that would justify noninvestment in cryptocurrencies throughout the sample period for different investment costs (in percent per year, applied to the absolute value of the weight in cryptocurrency). Panel D reports the inverse - for each of the priors in question, how high would costs have to be to result in non-investment up to that point? panel E reports investors' certainty equivalent return (CER) gains for the ex-post distribution of portfolio returns relative to the ex-post distribution of the market portfolio over the same period. Investors are assumed to have a relative risk aversion of 3. The sample consists of 106 monthly returns from May 2013 to February 2022.

Panel A: Optimal Portfolio Weights

Prior	Average	S.t.d	Lowest	Highest	Final	Fraction of positive	Fraction of above 0.5%	Fraction of above 1%	Fraction of above 2%	Fraction of above 5%	First Date Weight is Positive	Mean of Leverage	S.t.d of Leverage
Flat	0.146	0.084	-0.037	0.588	0.109	0.991	0.991	0.991	0.991	0.962	2013-06	1.082	0.082
2	0.053	0.010	0.022	0.066	0.063	1.000	1.000	1.000	1.000	0.613	all above	0.989	0.036
1	0.047	0.011	0.013	0.061	0.058	1.000	1.000	1.000	0.943	0.538	all above	0.979	0.035
0	0.041	0.012	0.004	0.056	0.053	1.000	0.991	0.953	0.943	0.217	all above	0.975	0.041
-1	0.034	0.013	-0.005	0.05	0.048	0.953	0.943	0.943	0.943	0.019	2013-10	0.968	0.045
-2	0.028	0.014	-0.014	0.045	0.043	0.943	1.000	0.991	0.726	0.000	2013-11	0.963	0.046
-5	0.009	0.016	-0.04	0.03	0.028	0.670	0.679	0.604	0.434	0.000	2013-11	0.943	0.047
-10	-0.022	0.021	-0.085	0.005	0.003	0.123	0.830	0.557	0.453	0.057	2021-02	0.910	0.051
-20	-0.084	0.031	-0.175	-0.046	-0.047	0.000	1.000	1.000	1.000	0.877	all below	0.853	0.055

Panel B: Ex-ante Certainty Equivalent Gains

Prior	2013	2014	2015	2016	2017	2018	2019	2020	2021	End of Sample
Flat	16.143	2.221	1.093	0.825	1.435	0.775	0.654	0.699	0.678	0.634
2	0.173	0.126	0.124	0.134	0.248	0.184	0.182	0.209	0.219	0.211
1	0.127	0.090	0.090	0.100	0.205	0.149	0.149	0.176	0.185	0.178
0	0.087	0.059	0.061	0.072	0.166	0.118	0.120	0.145	0.155	0.149
-1	0.055	0.036	0.038	0.048	0.131	0.091	0.094	0.116	0.128	0.122
-2	0.030	0.018	0.021	0.029	0.101	0.067	0.071	0.091	0.101	0.098
-5	0.000	0.001	0.000	0.001	0.034	0.017	0.021	0.033	0.043	0.042
-10	0.093	0.096	0.072	0.046	0.002	0.007	0.003	0.000	0.001	0.001
-20	0.808	0.737	0.609	0.478	0.242	0.257	0.208	0.157	0.118	0.118

Panel C: Investment Costs and Beliefs Required for Non-Investment

Cost/Time	2013	2014	2015	2016	2017	2018	2019	2020	2021	End of Sample
10 (min)	nan	nan	nan	nan	nan	nan	nan	nan	nan	nan
(max)	nan	nan	nan	nan	nan	nan	nan	nan	nan	nan
15 (min)	-2.90	-2.90	-2.90	-2.90	nan	nan	nan	nan	nan	nan
(max)	-2.60	-2.60	-2.60	-2.60	nan	nan	nan	nan	nan	nan
20 (min)	-3.80	-3.80	-3.80	-3.80	nan	nan	nan	nan	nan	nan
(max)	-1.70	-1.70	-1.70	-1.70	nan	nan	nan	nan	nan	nan
30 (min)	-5.50	-5.50	-5.50	-5.50	-5.50	-5.50	-5.50	-5.50	-5.50	-5.50
(max)	0.10	0.10	0.10	0.10	-1.70	-1.70	-1.70	-1.70	-1.80	-1.80
50 (min)	-8.90	-8.90	-8.90	-8.90	-8.90	-8.90	-8.90	-8.90	-8.90	-8.90
(max)	3.60	3.60	3.60	3.60	3.20	3.20	3.20	3.20	3.20	3.20
Lowest Preventing Cost	13.80	13.80	13.80	13.80	20.90	20.90	20.90	20.90	22.00	22.00
Corresponding Prior	-2.80	-2.80	-2.80	-2.80	-4.00	-4.00	-4.00	-4.00	-4.10	-4.10

Panel D: Minimum Cost for Non-Investment (in Percent)

Prior	2013	2014	2015	2016	2017	2018	2019	2020	2021	End of Sample
2	40.64	40.64	40.64	40.64	45.08	45.08	45.08	45.08	45.08	45.08
1	35.05	35.05	35.05	35.05	41.00	41.00	41.00	41.00	41.00	41.00
0	29.42	29.42	29.42	29.42	36.83	36.83	36.83	36.83	36.83	36.83
-1	23.80	23.80	23.80	23.80	32.76	32.76	32.76	32.76	32.76	32.76
-2	18.19	18.19	18.19	18.19	28.74	28.74	28.74	28.74	28.97	28.97
-5	1.00	1.00	1.00	2.31	16.63	16.63	16.63	16.63	19.07	19.07
-10	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	3.25	3.25
-20	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00

Panel E: Ex-post Certainty Equivalent Gains

Prior	2013	2014	2015	2016	2017	2018	2019	2020	2021	End of Sample
Flat	-1.633	-2.288	-1.386	-0.763	0.197	-0.171	-0.053	0.078	0.178	0.15
2	1.519	0.347	0.241	0.266	0.579	0.338	0.334	0.362	0.381	0.362
1	1.061	0.219	0.161	0.195	0.483	0.274	0.276	0.304	0.325	0.308
0	0.539	0.065	0.063	0.111	0.378	0.202	0.209	0.24	0.264	0.249
-1	-0.032	-0.111	-0.049	0.017	0.263	0.122	0.137	0.171	0.197	0.184
-2	-0.67	-0.317	-0.179	-0.089	0.138	0.032	0.056	0.093	0.123	0.113
-5	-2.953	-1.089	-0.666	-0.482	-0.297	-0.285	-0.229	-0.174	-0.129	-0.131
-10	-7.996	-2.902	-1.809	-1.375	-1.219	-0.983	-0.848	-0.747	-0.662	-0.649
-20	-22.642	-8.475	-5.317	-4.052	-3.798	-3.003	-2.624	-2.358	-2.143	-2.092
Max Gain	2.933	0.583	0.393	0.438	0.981	0.543	0.54	0.587	0.624	0.587
Optimal Prior	0.085	0.061	0.062	0.074	0.111	0.088	0.092	0.104	0.111	0.111
Positive CER, max prior	0.183	0.127	0.13	0.157	0.247	0.201	0.216	0.238	0.259	0.254
Positive CER, min prior	-0.009	-0.003	-0.005	-0.011	-0.03	-0.023	-0.026	-0.031	-0.035	-0.034

Table A3: Ex-ante Sharp Ratio Gains

This table reports estimates of the perceived monthly gain in Sharpe Ratios, in percentage points, from adding cryptocurrency to investors' existing portfolios. Panel A is for Bitcoin, Panel B is for an equal-weighted portfolio of cryptocurrency, and Panel C is for a value-weighted portfolio of cryptocurrency The reported values equal the difference between the Sharpe ratio of the baseline market portfolio that excludes cryptocurrency and the Sharpe ratio of the optimal portfolio that combines the market portfolio and cryptocurrency, assuming that investors have a constant relative risk aversion of 3. Years correspond to the end of the calendar year in question. Each row corresponds to a different prior belief, and each column corresponds to a specific end-of-year date when the investment decision is made. The sample consists of 106 monthly returns from May 2013 to February 2022.

Panel A: Ex-ante Sharpe Ratio Gains for BTC in Percentage

Prior	2013	2014	2015	2016	2017	2018	2019	2020	2021	End of Sample
Flat	88.009	26.423	16.980	13.930	18.907	12.867	11.776	12.671	11.106	10.339
2	3.674	2.675	2.877	3.009	5.035	3.953	4.105	4.404	4.440	4.304
1	2.714	1.937	2.111	2.338	4.219	3.210	3.469	3.903	3.743	3.706
0	1.910	1.317	1.494	1.687	3.507	2.557	2.852	3.274	3.182	3.241
-1	1.227	0.773	0.952	1.182	2.838	1.954	2.312	2.715	2.614	2.637
-2	0.658	0.371	0.506	0.720	2.143	1.490	1.768	2.199	2.143	2.158
-5	0.011	0.057	0.007	0.014	0.636	0.330	0.551	0.852	0.913	0.848
-10	2.437	2.478	1.835	1.249	0.130	0.283	0.102	0.003	0.004	0.002
-20	14.044	13.239	11.442	9.721	5.876	6.337	4.831	3.332	2.871	2.880

Panel B: Ex-ante Sharpe Ratio Gains for EW in Percentage

Prior	2013	2014	2015	2016	2017	2018	2019	2020	2021	End of Sample
Flat	86.858	27.958	16.182	12.173	18.269	12.911	10.277	9.521	9.875	9.276
2	2.681	2.141	1.965	2.038	3.984	3.288	2.995	3.035	3.274	3.321
1	2.271	1.807	1.654	1.736	3.662	2.971	2.725	2.729	2.929	3.029
0	1.880	1.450	1.379	1.452	3.340	2.655	2.453	2.386	2.676	2.695
-1	1.528	1.164	1.101	1.201	2.997	2.382	2.212	2.134	2.508	2.428
-2	1.202	0.901	0.863	0.981	2.694	2.106	1.978	1.911	2.269	2.158
-5	0.443	0.306	0.305	0.394	1.769	1.365	1.222	1.253	1.569	1.473
-10	0.006	0.016	0.010	0.000	0.613	0.452	0.401	0.442	0.633	0.635
-20	2.277	2.110	1.786	1.333	0.111	0.167	0.140	0.067	0.003	0.005

Panel C: Ex-ante Sharpe Ratio Gains for VW in Percentage

Prior	2013	2014	2015	2016	2017	2018	2019	2020	2021	End of Sample
Flat	86.676	26.154	16.027	12.772	19.154	12.159	10.818	11.051	10.476	10.073
2	3.567	2.619	2.585	2.764	4.976	3.865	3.715	4.101	4.091	4.096
1	2.680	1.910	1.939	2.125	4.237	3.235	3.113	3.480	3.525	3.531
0	1.880	1.286	1.343	1.525	3.454	2.672	2.540	2.984	3.028	3.036
-1	1.212	0.801	0.872	1.034	2.809	2.112	2.070	2.448	2.558	2.541
-2	0.691	0.405	0.482	0.643	2.182	1.574	1.592	1.949	2.040	2.038
-5	0.001	0.024	0.002	0.015	0.752	0.409	0.509	0.704	0.873	0.901
-10	2.031	2.070	1.596	1.030	0.048	0.166	0.084	0.003	0.020	0.013
-20	12.575	11.621	10.323	8.576	4.778	5.321	4.306	3.017	2.367	2.379

Table A4: Investment Costs and Certainty Equivalent Return Gains

This table considers the effect of annual investment costs on certainty equivalent return (CER) gains from cryptocurrency. Panel A reports the ex-ante monthly CER gain, in percentage points, from adding cryptocurrency to investors' existing portfolios, which corresponds to the end of the sample period. Panel B reports analogous estimates of ex-post monthly CER gains. Each row corresponds to a different prior belief. Each column corresponds to a different combination of a cryptocurrency asset and an annual investment cost (in percent per year, applied to the absolute value of the weight in cryptocurrency). Investors are assumed to have a relative risk aversion of 3. The sample consists of 106 monthly returns from May 2013 to February 2022.

Panel A: Ex-ante Gains

	Annual cost 20%	Annual cost 10%	Annual cost 5%	Annual cost 2%
BTC				
Flat	0.479	0.579	0.629	0.660
2	0.112	0.171	0.200	0.218
1	0.086	0.140	0.166	0.183
0	0.063	0.112	0.137	0.151
-1	0.042	0.086	0.108	0.121
-2	0.024	0.064	0.083	0.095
-5	-0.008	0.016	0.029	0.036
-10	-0.002	-0.001	-0.001	0.000
-20	0.051	0.097	0.119	0.133
EW				
Flat	0.464	0.508	0.530	0.544
2	0.113	0.137	0.149	0.157
1	0.101	0.124	0.135	0.142
0	0.088	0.109	0.120	0.127
-1	0.076	0.097	0.107	0.113
-2	0.066	0.086	0.095	0.101
-5	0.039	0.055	0.063	0.068
-10	0.009	0.020	0.025	0.028
-20	-0.002	-0.001	0.000	0.000
VW				
Flat	0.452	0.543	0.588	0.616
2	0.106	0.158	0.185	0.201
1	0.082	0.130	0.154	0.169
0	0.060	0.105	0.127	0.140
-1	0.042	0.082	0.102	0.114
-2	0.027	0.063	0.080	0.091
-5	-0.005	0.019	0.031	0.038
-10	-0.005	-0.002	-0.001	0.000
-20	0.040	0.079	0.099	0.111

Panel B: Ex-post Gains

	Annual cost 20%	Annual cost 10%	Annual cost 5%	Annual cost 2%
BTC				
Flat	-0.072	0.063	0.131	0.172
2	0.306	0.355	0.38	0.395
1	0.258	0.301	0.323	0.336
0	0.203	0.240	0.259	0.270
-1	0.142	0.174	0.190	0.199
-2	0.074	0.100	0.114	0.122
-5	-0.174	-0.161	-0.154	-0.15
-10	-0.771	-0.749	-0.738	-0.731
-20	-2.512	-2.433	-2.393	-2.369
EW				
Flat	0.283	0.347	0.379	0.398
2	0.248	0.269	0.279	0.285
1	0.221	0.24	0.25	0.256
0	0.192	0.209	0.218	0.224
-1	0.161	0.177	0.185	0.190
-2	0.128	0.143	0.151	0.155
-5	0.019	0.031	0.036	0.040
-10	-0.194	-0.188	-0.185	-0.183
-20	-0.768	-0.757	-0.752	-0.749
VW				
Flat	-0.097	0.027	0.089	0.126
2	0.272	0.317	0.339	0.353
1	0.229	0.269	0.288	0.300
0	0.181	0.215	0.232	0.242
-1	0.126	0.155	0.17	0.178
-2	0.064	0.088	0.101	0.108
-5	-0.155	-0.143	-0.137	-0.133
-10	-0.686	-0.667	-0.658	-0.652
-20	-2.231	-2.162	-2.127	-2.106

Table A5: Snapshot and Cumulative Cutoff Priors for the Equal-Weighted and Value-Weighted Cryptocurrency Portfolios, with ambiguity aversion $\tau = 4$

This table reports the snapshot and accumulative cutoff priors for non-investment in the equal-weighted and Value-weighted cryptocurrency portfolio in a robust portfolio choice framework.

Panel A: Snapshot Cutoff Priors

	2013	2014	2015	2016	2017	2018	2019	2020	2021	End of Sample
ew-rf	-0.095	-0.091	-0.093	-0.102	-0.17	-0.163	-0.164	-0.172	-0.196	-0.193
vw-rf	-0.049	-0.044	-0.049	-0.056	-0.091	-0.081	-0.086	-0.098	-0.109	-0.107

Panel B: Cumulative Cutoff Priors

	2013	2014	2015	2016	2017	2018	2019	2020	2021	End of Sample
ew-rf	-0.097	-0.097	-0.097	-0.102	-0.170	-0.175	-0.175	-0.175	-0.197	-0.197
vw-rf	-0.052	-0.052	-0.052	-0.056	-0.091	-0.091	-0.091	-0.098	-0.111	-0.111

Table A6: Optimal Portfolio Weights for the Equal-Weighted and Value-Weighted Cryptocurrency Portfolios, with ambiguity aversion $\tau = 4$

This table reports cryptocurrency portfolio weights for different prior beliefs about the average monthly cryptocurrency excess return at the beginning of the sample period, with ambiguity aversion $\tau = 4$. Panel A considers an equal-weighted cryptocurrency portfolio, and Panel B considers a value-weighted cryptocurrency portfolio. Each row corresponds to a different prior belief. For each prior, the columns indicate a range of attributes of the distribution of weights over the sample period - the average, lowest, highest, final (i.e. end-of-sample) weight, the fraction of months that are above zero, the fraction of weights whose absolute value exceeds 0.5%, 1.5%, 2.5% and 5%, the first date in the sample when the weight is positive, as well as the mean and standard deviation of leverage choices. The sample consists of 106 monthly returns from May 2013 to February 2022.

Panel A: Equal-Weighted Cryptocurrency Portfolio

Prior	Average	S.t.d	Lowest	Highest	Final	Fraction positive	Fraction above 0.5%	Fraction above 1%	Fraction above 2%	Fraction above 5%	First Date Weight is Positive	Mean of Leverage	S.t.d of Leverage
2	0.010	0.002	0.003	0.013	0.012	1.000	0.943	0.585	0.000	0.000	All above	0.411	0.013
1	0.010	0.002	0.002	0.012	0.012	1.000	0.943	0.547	0.000	0.000	All above	0.41	0.013
0	0.009	0.002	0.001	0.012	0.011	1.000	0.943	0.519	0.000	0.000	All above	0.409	0.013
-1	0.008	0.003	0.000	0.011	0.010	1.000	0.943	0.292	0.000	0.000	All above	0.407	0.012
-2	0.007	0.003	-0.001	0.010	0.010	0.943	0.943	0.113	0.000	0.000	2013-Nov	0.406	0.011
-5	0.005	0.003	-0.004	0.008	0.008	0.943	0.557	0.000	0.000	0.000	2014-Nov	0.404	0.011
-10	0.002	0.004	-0.009	0.006	0.005	0.632	0.179	0.000	0.000	0.000	2016-Jun	0.398	0.014
-20	-0.005	0.005	-0.019	0.000	0.000	0.000	0.453	0.057	0.000	0.000	All Below	0.396	0.016

Panel B: Value-Weighted Cryptocurrency Portfolio

Prior	Average	S.t.d	Lowest	Highest	Final	Fraction positive	Fraction above 0.5%	Fraction above 1%	Fraction above 2%	Fraction above 5%	First Date Weight is Positive	Mean of Leverage	S.t.d of Leverage
2	0.023	0.004	0.009	0.028	0.027	1.000	1.000	0.981	0.736	0.000	All above	0.419	0.014
1	0.020	0.005	0.006	0.026	0.025	1.000	1.000	0.943	0.557	0.000	All above	0.416	0.015
0	0.017	0.005	0.002	0.024	0.023	1.000	0.943	0.943	0.425	0.000	All above	0.414	0.015
-1	0.015	0.005	-0.002	0.021	0.020	0.953	0.943	0.868	0.142	0.000	2013-Oct	0.411	0.018
-2	0.012	0.006	-0.006	0.019	0.018	0.943	0.972	0.575	0.000	0.000	2013-Nov	0.409	0.019
-5	0.004	0.007	-0.017	0.013	0.012	0.670	0.604	0.217	0.000	0.000	2014-Nov	0.402	0.020
-10	-0.009	0.009	-0.036	0.002	0.001	0.132	0.509	0.453	0.057	0.000	2021-Jan	0.387	0.023
-20	-0.035	0.013	-0.072	-0.019	-0.020	0.000	1.000	1.000	0.934	0.057	All Below	0.357	0.026

Table A7: Optimal Weights of the Equal-Weighted and Value-Weighted Cryptocurrency Portfolios for Different Volatility Priors

This table explores how different prior beliefs about volatility affect the average and end-of-sample weights on the equal-weighted and value-weighted cryptocurrency Portfolios across different prior means

Panel A: Average Weights of the Equal-Weighted Cryptocurrency Portrolio

Volatility \EW prior	2	1	0	-1	-2	-5	-10	-20
0.2x Sample	0.147	0.133	0.12	0.107	0.093	0.054	-0.012	-0.139
0.5x Sample	0.068	0.063	0.058	0.053	0.048	0.032	0.007	-0.043
Baseline	0.024	0.023	0.021	0.019	0.018	0.013	0.004	-0.013
2x Sample	0.007	0.007	0.006	0.006	0.005	0.004	0.002	-0.003
5x Sample	0.002	0.001	0.001	0.001	0.001	0.001	0.001	0.000

Panel B: End-of-Sample Weights of the Equal-Weighted Cryptocurrency Portrolio

Volatility \EW prior	2	1	0	-1	-2	-5	-10	-20
0.2x Sample	0.160	0.152	0.143	0.136	0.127	0.103	0.063	-0.014
0.5x Sample	0.079	0.075	0.071	0.067	0.063	0.052	0.032	-0.006
Baseline	0.029	0.027	0.026	0.025	0.023	0.019	0.012	-0.001
2x Sample	0.009	0.008	0.008	0.008	0.007	0.006	0.004	0.000
5x Sample	0.002	0.002	0.002	0.002	0.001	0.001	0.001	0.000

Panel C: Average Weights of the Value-Weighted Cryptocurrency Portrolio

Volatility \VW prior	2	1	0	-1	-2	-5	-10	-20
0.2x Sample	0.340	0.290	0.239	0.189	0.139	-0.009	-0.249	-0.699
0.5x Sample	0.151	0.132	0.113	0.094	0.075	0.018	-0.076	-0.257
Baseline	0.053	0.047	0.041	0.034	0.028	0.009	-0.022	-0.084
2x Sample	0.016	0.014	0.012	0.011	0.009	0.004	-0.005	-0.022
5x Sample	0.003	0.003	0.003	0.002	0.002	0.001	0.000	-0.003

Panel D: End-of-Sample Weights of the Value-Weighted Cryptocurrency Portrolio

Volatility \VW prior	2	1	0	-1	-2	-5	-10	-20
0.2x Sample	0.351	0.322	0.292	0.262	0.232	0.144	0.001	-0.257
0.5x Sample	0.175	0.161	0.146	0.131	0.117	0.074	0.004	-0.129
Baseline	0.063	0.058	0.053	0.048	0.043	0.028	0.003	-0.047
2x Sample	0.019	0.018	0.016	0.015	0.014	0.009	0.002	-0.012
5x Sample	0.004	0.004	0.003	0.003	0.003	0.002	0.001	-0.001