The Cryptocurrency Participation Puzzle

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

First Draft: August 2022

Abstract

We show that ongoing zero portfolio weights in cryptocurrency are surprisingly difficult to generate in a standard Bayesian portfolio theory framework. With ten years of prior data, equity market investors would need very pessimistic priors on mean returns to justify never having bought cryptocurrency: -10.6% per month for Bitcoin, and -19.6% per month for a diversified portfolio of cryptocurrencies. Moreover, most priors that involve never purchasing cryptocurrency imply that investors should *short* cryptocurrency. Optimal absolute weights are generally small but non-trivial (1-5%), frequently positive, and fairly smooth despite returns being volatile. Under a wide range of priors, the certainty equivalent gains from cryptocurrency are comparable to international diversification and exceed the size anomaly. Costs (ambiguity aversion, storage, fees) would need to be enormous to justify never trading, over 21% per year for Bitcoin and 39% for a diversified cryptocurrency portfolio.

*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, and the MSUFCU conference for helpful comments and suggestions. All remaining errors are our own.

Stanislaw Ulam: *"I define a coward as someone who will not bet when you offer him two-to-one odds and let him choose his side."*
Paul Samuelson: *"You mean will not make a sufficiently small bet."*

(Recounted in Samuelson 1963)

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 the Consumption Capital Asset Pricing Model (CCAPM) and other models, given the risk-adjusted expected returns from holding cryptocurrency. The avoidance of cryptocurrency investments is particularly prevalent among those less well-off. Yet liquidity constraints are not the only reason for not holding cryptocurrencies – many individuals who have substantial liquid assets also hold no cryptocurrencies. Insights into the causes of non-participation 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".[1] 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 two 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. An individual investor may obtain cash flows by selling to somebody else, but this does not explain why the asset class as a whole attained non-zero prices in the first place, or why it has them today.

While opinions change over time, many market observers and sophisticated investors view cryptocurrencies as a bubble or a Ponzi scheme (e.g. Griffin and Shams 2020).[2] Cryptocurrency skeptics often cite such views to justify holding a cryptocurrency portfolio weight of zero since the invention of Bitcoin in 2009.

[1] Respectively, Kuhnen and Miu 2017, Bogan 2008, Kuhnen and Miu 2017, Mankiw and Zeldes 1991, and Briggs et al. 2021. The modified sentences reflect updated information about individual cryptocurrency participation rates from https://www.finder.com/how-many-people-own-cryptocurrency. Individual equity non-participation rates are lower - around 50% of households do not hold equities (Campbell 2006, Calvet, Campbell, and Sodini 2007).

[2] For example, in an interview to CNBC on December 12, 2021, Rich Bernstein said that "Cryptos are the biggest financial bubble ever in history."

The view that cryptocurrencies are worthless because they cannot produce cash flows is based on venerable discounted cash flow models. However, even if one believes that cryptocurrencies *ought* to be priced at zero, this does not necessarily mean they will end up there. The positive aspect of discounted cash flow models can be represented as *priors* about the distribution of expected returns. In this framework, saying that Bitcoin is a scam or a bubble implies that it has negative expected returns over some future horizon, which we dub "pessimistic priors". While such priors may be based on compelling economic logic, they are still priors nonetheless, and should be amenable to quantification and updating via Bayes' rule.

We use a Bayesian portfolio theory framework (Pástor 2000, Tu and Zhou 2010) to address the elephant-in-the-room question about cryptocurrency, which the literature (to our knowledge) has largely, and surprisingly, not tackled head on. That is, *should investors actually buy any*? We highlight two important attributes of our analyses. First, Bayesian portfolio theory is especially suited to studying cryptocurrency participation, where differences in beliefs are large and first-order. Second, our analysis is normative. We take the perspective of a relatively sophisticated and unconstrained investor with access to different financial markets (such as futures exchanges), and vary his prior beliefs about cryptocurrency. Our aim is to quantify how easily pessimistic priors can justify a zero allocation to cryptocurrencies.[3] Unlike the equity participation literature, which compares purchasing equities to zero weights, we evaluate avoiding positive *or* negative weights.

Our main finding is that an allocation of precisely zero to cryptocurrencies, 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. And 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.

In our baseline 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 investors' level of certainty in these priors, we assume that they have observed ten equivalent years of data before the series began. We assume that investors initially hold the US market portfolio, and choose an optimal

[3]Several papers have documented high cryptocurrency returns. While high returns may imply that investment is sensible, prior studies have not explicitly tackled the normative implications or quantified the role of prior beliefs. We discuss these papers in the "Related Literature" section.

portfolio that combines positions in cryptocurrency, subject to reasonable levels of risk aversion. In subsequent analyses, we consider different beliefs about volatility and correlations, different levels of certainty about prior beliefs, and alternative initial portfolios.

We find that investors' prior beliefs about mean Bitcoin returns would need to be lower than -10.3% per month to justify non-investment at the end of the sample period (February 2022), and lower than -10.6% to justify never investing during the sample period. For the equally-weighted cryptocurrency portfolio, the results are even more extreme: priors would need to be lower than -19.2% for end-of-sample non-investment, and -19.6% for never investing. These estimates, however, understate the puzzle since large negative priors 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. Of priors that generate zero weights if we restrict short-selling, nearly all entail nontrivial negative weights for unconstrained investors. This reflects an old result from portfolio theory that is not unique to cryptocurrency – an optimal portfolio will generally have some weight, positive or negative, on all non-redundant assets.

Investors' optimal weights have three important properties - they are generally i) small, ii) smooth, and iii) frequently positive. Across a wide range of priors about mean cryptocurrency returns (between 2% and -20% per month), optimal weights range from a high of 7.3% to a low of -19.8%, with most absolute weights ranging between 1% and 5%. Because investors are well-calibrated about cryptocurrency's high volatility, they only take small positions. Despite this volatility, the optimal weights move smoothly over time, and (other than diffuse priors) do not exhibit large swings. With ten years of data before the series starts, and more data added over time, the impact of any later price movement is relatively small. As such, our conclusions are unlikely to change by adding new data. Close-to-zero *average* weights tend to correspond to negative weights early in the sample and positive weights later on, rather than consistent zero weights. While pessimistic investors may have taken zero weights due to short sales constraints early in the sample, this is less compelling after 2017 when futures markets became available.

Despite being small, desired weights are nontrivial. The absolute value of optimal weights in Bitcoin always exceeds 0.9% at some point in the sample for *any* ten-year prior (even including early short sales constraints), and the desired weights absent short sales constraints would be even larger, exceeding 3.6%. For equal-weighted cryptocurrency, short sales constraints result in zero weights for sufficiently pessimistic priors, but desired absolute weights always exceed 1.6% at some point for any ten-year prior. In other words, non-investment is not easily explained by weights being too small to be "worth it".

Next, we consider how non-participation in cryptocurrency markets compares to other asset markets, such as the "home bias puzzle" in international equities (French and Poterba 1991, Tesar and Werner 1995). Cryptocurrencies have two important conceptual differences compared to other assets where the average investor has zero weights.[4] First, investors who do not purchase assets such as international equities or convertible bonds are anecdotally much less likely to justify non-participation on the basis that these assets are a bubble that is going to zero. Second, the normative advice given by academics in such cases is generally 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. This attitude does not seem to be widely applied to cryptocurrency.

One possible reason for the discrepancy in investment advice across assets could be differences in the diversification benefits they provide. We show that, under a wide range of priors, an equity market investor would perceive Bitcoin as generating sizable certainty equivalent of returns (CER) gains at the end of the sample. Priors between -2% and 2% per month produced perceived portfolio-level gains of 10 to 23 b.p. per month. These magnitudes exceed the CER gains on the MSCI world (ex-US) and SMB portfolios. The size of the benefits from adding cryptocurrency is thus at least comparable to other assets, even for various initially skeptical priors and relatively small weights.

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 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 Bitcoin. For the equally weighted cryptocurrency portfolio, costs would need to exceed 30% per year. In other words, to deter

[4]The stock market participation literature assumes that the choice set, and thus the puzzle, is between positive weights in equities and the risk-free rate. One could equally ask about the convertible bond participation puzzle, the mortgage-backed securities puzzle, and dozens of similar assets that are even more likely to be allocated zero weights.

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 quite small.

We also investigate investors' ex-post satisfaction with their cryptocurrency investments in two ways. First, given their chosen portfolio, are they happy with the ex-post realized return distribution if it were to continue? Second, given their beliefs at the time, would they have continued to invest in cryptocurrency? Note that these are separate questions - one could continue with a trade (thinking it will produce good returns given one's priors) despite being unhappy with the returns up to that point. Unsurprisingly, a wide range of optimistic priors lead to ex-post happiness, although at extreme levels of optimism (above 25.1% mean excess returns), investors are ex-post unhappy due to the large volatility that their beliefs expose them to. More strikingly, even initially pessimistic priors above -3.4% eventually lead to long positions and ex-post happiness, whereas more pessimistic priors lead to ex-post unhappiness. While we are agnostic over priors, these results show that acting on highly pessimistic beliefs was ex-post unsatisfactory.

Zero weights are not impossible. They arise most easily from posterior beliefs of excess returns close to zero, where high volatility makes it undesirable to either be long *or* short. Consistent zero weights require greater dogmatism (e.g. assuming one has seen 30 or 50 years of data) *and* priors that cryptocurrency will earn returns somewhat lower than the risk-free rate. Greater dogmatism will act as an anchor that weakens the impact of the high and volatile realized returns, and prevents the posterior changing much. Mildly pessimistic priors, if held quite dogmatically, can combine with large positive sample returns to give posterior means consistently close to zero, and thus weights also consistently near zero. If investors are presumed to have seen 50 years of past data, persistent non-investment arises for priors of -2.4% for Bitcoin and -4.6% for the equally weighted cryptocurrency portfolio. Intuitively, these priors map to strong beliefs that Bitcoin is a slowly deflating bubble, rather than a bubble about to burst. In the limit, sufficiently large dogmatism practically amounts to a formalized license to ignore the observed returns.

We extend our baseline specification in different ways. Our conclusions remain robust if investors start with different equity portfolios. Prior beliefs that cryptocurrency is positively correlated with equity markets (instead of uncorrelated in the baseline case) have only moderate effects on cutoff priors. Beliefs in higher volatility shrink desired weights but do not significantly affect cutoff beliefs for exactly zero weights. Similar effects are produced by model uncertainty, such as robust decision-making (Hansen and Sargent

2011). These analyses are not meant to be an exhaustive range of possible beliefs, but rather a reasonable formalization of skepticism about cryptocurrency. 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 reduce this to a prior over the current month's returns and reevaluate, a similar analysis ought to apply. We discuss these and other possible extensions at the end of the paper. While our tools are assuredly not the final word on determining optimal portfolio allocations, they are reasonably standard, and a good framework for rigorously incorporating different beliefs without relying on ad hoc justifications and informal rhetorical arguments. Finally, in the Appendix, we describe some reasons why investors might *not* necessarily have 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 in the absence of a fully fleshed out model.

Overall, we find that the case for simultaneously believing that cryptocurrency is a bubble and also taking weights of zero to be not nearly as straightforward as is commonly assumed. While Bayesian methods can be complicated to implement, their logic is quite intuitive. Large positive returns ought to lead many initially pessimistic investors to eventually become optimistic. Portfolio theory generally implies non-zero weights in all non-redundant assets. Strong beliefs that cryptocurrency is a bubble with large negative returns imply one should short cryptocurrency, rather than not participate at all. 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 (Hartzmark and Solomon 2021) is a license for skepticism and non-participation, 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

Our paper adds to research on cryptocurrencies. One strand of the literature studies cryptocurrency returns. Y. Liu, Tsyvinski, and Wu 2022 provide evidence 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 characterize the investable universe of cryptocurrencies and its composition, realized performance, volatility, and correlation with traditional assets. Yi, Xu, and G.-J. Wang 2018, Zhang et al. 2018, and Iyer 2022 describe 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 studying the set of beliefs that could justify non-participation in the cryptocurrency market, given the realized returns. Relative to existing papers, we focus explicitly on the role of beliefs. Rather than using sample returns to dispute the basis for pessimistic priors, we instead 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 provide evidence that the prices of Bitcoin and other cryptocurrencies were inflated by the supply of Tether during the 2017 boom. Makarov and Schoar 2019 study price discovery in cryptocurrency markets and its relation to market segmentation and investors' 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. Much of this literature studies explanations for households' non-investment in equities (e.g., Mankiw and Zeldes 1991, Campbell 2006, and Calvet, Campbell, and Sodini 2007, and many others). Unlike our current work, this literature is nearly all positive in nature, trying to explain the biases or costs that might drive the choices in practice. Instead, we focus on the normative aspect that mostly is just assumed in the case of equities, namely whether investors *should* be holding the assets. Indeed, there is no shortage of people who take the view for cryptocurrency that the puzzle to be explained is the opposite one, namely why anyone *is* buying any at all. We are interested in establishing a normative baseline under reasonable estimates of costs and barriers to trade, and leave to future research the important question of why, in practice, 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 a much broader range of assets than just equities, and potential explanations for equities ought to 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 literature on equities), but failing to either buy *or short* it, if the latter option 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 of these problems are exacerbated for small coins, whose returns are extremely volatile.

To balance this trade-off, we utilize screens that exclude obvious data errors as well as highly illiquid coins. This allows us to avoid altering the reported return data with arbitrary changes and cutoffs, like winsorizing. The screens we use are:

-Negative prices and market capitalizations are set to missing.

-Returns are dropped if the changes in market capitalization and price that day seem incongruent. That is, the absolute value of percentage price change minus percentage market capitalization change is greater than two.[5]

-Returns are dropped if 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.

-Returns are dropped if the turnover the previous day (total volume divided by market capitalization) is less than 0.1%, or greater than 200%.

-Returns are dropped if the market capitalization does not exceed $100 million on at least three of the previous 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.[6] 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.

[5]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.

[6]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.

Factor portfolios for SMB, HML and UMD are taken from Ken French's website. MSCI World Ex-US returns are taken from SP 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. We assume that investors have priors over means, standard deviations, and correlations, and solve a portfolio choice problem. The investor has access to leverage, and so is choosing the optimal portfolio as a combination of the risk-free asset and the estimated tangency portfolio among the set of risky assets, given his priors and the observed data.

We assume that there are N risky assets and a risk-free asset. For most of the paper, the final portfolio choice is between excess market returns and cryptocurrencies. In later analyses we also use other portfolios (e.g., the size portfolio), so we calculate properties initially for all N risky assets. 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

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

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

where γ is the relative risk-aversion coefficient, equal to three. The level of risk aversion does not affect the estimated tangency portfolio or the investor's 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. It is well known (Markowitz 1952) that 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.

To calculate the optimal weights w^* in equation 3, the true values of μ and V are needed. Since these are not actually known, the usual technique is to estimate sample versions from observed data, and treat them as the true values. However, this gives rise to a parameter uncertainty problem, and the utility associated with weights calculated from sample parameters can be substantially different from $U(w^*)$ under the true parameters.

We instead use a Bayesian method for optimal portfolio choice. The Bayesian optimal portfolio is obtained by maximizing the expected utility under the predictive distribution.

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

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

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.

In comparing equation (6) with equation (1), expected utility is maximized in the Bayesian and classic framework under the predictive and true distributions, respectively. However, the evaluation of equation (1) requires treating the two-step estimates as the true parameters and is hence subject to estimation error, while the Bayesian approach accounts for the estimation error automatically. “Predictive density" underlies the fact that investors update their beliefs on the distribution of the parameters via Bayesian methods, and then update the asset return distributions according to the updated parameter posterior distributions. With different data realizations or different prior beliefs, investors can infer a different predictive density and form different optimal portfolios.

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,

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

An inspection of 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 effect 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,

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

Notice that this expected utility is evaluated based on the same μ^* and V^* from the predictive density. We can define the difference

$$EU_{Bayes} - EU_{res} \tag{11}$$

as the ex-ante or perceived certainty-equivalent return (CER) gain from cryptocurrency.

Similarly, we can measure the ex-post gain from investing in cryptocurrency for a given 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, we define the ex-post certainty-equivalent return (CER) as

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

The difference between these two terms (U_{Bayes} - U_{res}) is then the ex-post certainty-equivalent return (CER) gain from investing in cryptocurrency. As any weight 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 in 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.

4 Results

4.1 Prior Beliefs and Cryptocurrency Participation

We begin by examining how pessimistic an investor's prior beliefs about cryptocurrency would need to be to justify non-participation at each point in 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 will map to a desire to buy the asset, and all lower posterior beliefs map to a desire to short the asset. The relevant priors are over excess returns, and throughout the paper statements about returns reference excess returns unless otherwise noted. Solving for these cutoff, zero-weight prior beliefs requires specifying 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.

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 the equal-weighted portfolio has higher mean returns than Bitcoin, prior beliefs would have needed to be even more pessimistic, ranging from -9.1% per month in 2014, to -19.2% per month at the end of the sample. Row 3 repeats the analysis for the value-weighted cryptocurrency portfolio. The 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. For brevity, we only report the results for Bitcoin and the equal-weighted cryptocurrency portfolio in the paper, and include the main results for the value-weighted portfolio in the Internet Appendix.

In Panel B, we consider a related question - what prior beliefs would lead investors to never take a positive weight prior to the annual snapshot in question. 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, studying local price minima and maxima gives similar results 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

The previous section considered the cutoff beliefs that would lead to positive weights. These are especially important when investors cannot short sell. In such cases, beliefs implying negative desired weights lead to zero weights. In terms of the ability to short sell cryptocurrencies, for Bitcoin a clear delineation can be made 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. After December 2017, our basic framework predicts that everyone whose beliefs were not *exactly equal* to the numbers in Table 2 should have taken either a long or a short position in Bitcoin. This reflects the fact that zero weights are generally not optimal for most assets under portfolio theory.

While the results in Table 2 demonstrate how difficult it is to justify zero weights year after year, it is also important to understand *how much* one should invest in cryptocurrency. If optimal weights are very close to zero, people may follow a heuristic where very small trades are considered "not worth it". Since cryptocurrency is very volatile and has had large returns, small weights may still have considerable effects, so the heuristic is not obviously a good one, despite its plausibility. We calculate cryptocurrency portfolio weights for a wide array of opinions about cryptocurrency, including mildly optimistic priors (2% and 1% per month), pessimistic priors (-1%, -2%, -5%, -10% and -20%), neutral (or no risk-exposure, market-efficiency) priors of 0% per month, and flat/diffuse priors where the investor believes in the sample mean observed up to that point. While one could also study highly optimistic investors, we suspect (anecdotally) that most optimists actually held cryptocurrency, and so do not have a participation puzzle.

Panel A of Table 3 summarizes the range of optimal Bitcoin portfolio weights, where each row corresponds to a different prior belief. Column 1 shows that average weights range from 5.9% for the 2% prior to -9.6% for the -20% prior. The largest absolute weights result from diffuse priors, where the average weight is 16.0% (though, as we will see later, these large weights do not lead to ex-post desirable portfolios). Columns 2 and 3 show maximum and minimum weights over the sample period. 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%).

[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/.

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 whose absolute values are above 0.5%, 1%, 2% and 5% respectively. If weights above 2% in either direction 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 only reached with any frequency 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, the desired weights in Bitcoin are not especially large in either direction (mostly in the 1-5% range), even under a very 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 will 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 equal-weighted cryptocurrency portfolio, and shows similar patterns to Panel A. 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 graphs the evolution of weights over time. Panel A examines the optimal weights for Bitcoin at each of the priors, and Panel B examines the equal-weighted cryptocurrency portfolio. These graphs show that, 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 in implementing 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, unless there is a prolonged period of very negative returns, enough to offset both the informative priors imposed and the existing sample data already incorporated. [8] The intuition for this result is that while cryptocurrency monthly returns are volatile, posteriors about means change relatively slowly once an informative prior is imposed, even more so additional data has been observed. As long as the volatility of returns is reasonably stable (and understood), the optimal actions do not change rapidly.

Panels C and D of Figure 2 examine how maximum desired weights over the sample period vary with priors. Specifically, for each prior we compute the maximum absolute weight that an investor desires over the sample (absent short sale constraints). We then calculate the minimum across all ten-year priors. 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 in the sample (with the minimum corresponding to a prior of -0.077). The shape of both graphs is a v-shape, implying that the 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 (without shorting before 2017), investors with all priors desire an absolute weight of at least 0.9% at some point in the sample (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

[8] Some results in the paper, like the "priors required to never invest", are a ratchet that can only get more negative over time, regardless of what future data arrives.

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 one that 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 CER gains from these investments to the CER gains from cryptocurrencies at the end of the sample period. 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 very wide range of priors (i.e. the "Bitcoin to MSCI" ratio is greater than the "SMB to MSCI" ratio), though not HML and UMD . For the equal-weighted cryptocurrency portfolio, 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 Ambiguity Aversion, Storage Costs and Other Frictions

Next, we examine the effect of costs that may deter investors from taking non-zero positions. One such cost is ambiguity aversion (Epstein and T. Wang 1994). Cryptocurrencies may create significant uncertainty over their value proposition, and new technology (e.g., Blockchain technology, ledger storage, public and private key cryptography) which can generate ambiguity/entry costs. 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 the assets during the life of the owner, is also technically challenging. Finally, there are more basic transaction costs, including bid-ask spreads, impact, and exchange fees.

While some of these are likely fixed costs (such as 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 (e.g. F. Gomes and Michaelides 2005, Haliassos and Michaelides 2003, Abel, Eberly, and Panageas 2013) for future research.[9] These costs a simple way to represent various reasons why investors may be reluctant to trade cryptocurrencies, and consequently choose a zero weight even when the costless version of portfolio theory would imply 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. Assigning a magnitude for some of these costs is not straightforward - obviously for psychological costs like ambiguity aversion, but even spreads and market impact early on 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 the equal-weighted cryptocurrency portfolio) take a measure of costs, and show the range of prior beliefs that correspond to non-investment in cryptocurrencies up to the 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 perform the reverse calculation - for each prior, we solve for the minimum cost that would justify non-investment for all annual snapshots up to each sample

[9] 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.

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 the cost estimates are. Costs that deter investment *start* at over 20% per year for Bitcoin and 39% for the equal-weighted cryptocurrency portfolio. The large estimates mitigate concerns about the exact form of the costs, as the costs required to deter investment are large to the point of implausibility.

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 in time we computed how investors would have perceived cryptocurrency on a forward-looking basis, given their prior beliefs. This is separate from whether the trades actually performed well or not. 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 over the sample period, ex-post, however, 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%, ranging from 0.41% at 2% priors to 0.13% for priors of -2%. That is to say, investors with ten year priors 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 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, we vary our main specification, by 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 is a reasonable versus unreasonable level of dogmatism about cryptocurrency returns. 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). We consider this in Table 8 Panel A, which examines the cutoff priors that map to zero investment up to given points in time.

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 Table 8 Panel B, we examine the effect of the investor holding 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. As such, we examine how prior beliefs about volatility affect the weights chosen for various prior 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 cryptocurrency portfolio 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 final 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.

Panel F examines cutoff beliefs if investors ignore early data 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.

Panels G and H examine a robust version of portfolio choice that allows for 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. After they make their choice, an adversarial agent pays a cost to perturb the probability distribution of returns in a way designed to be maximally costly for the investor's 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. This max-min or min-max optimization problem is conceptually similar to Generative Adversarial Networks in the Computer Science literature (Goodfellow, Shlens, and Szegedy 2014).[10] Details are provided in the internet appendix. This model uncertainty ends up operating similar to belief in higher volatility. 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).

Overall, varying initial assets makes little difference to our conclusions. Model uncertainty and beliefs in higher volatility shrink desired weights, but do not directly produce zero weights without additional frictions. Beliefs about higher correlations or greater dogmatism have more substantial effects on cutoff beliefs required for zero investment.

5 Alternative Modeling Choices

Our analysis illustrates reasonable ways to think about cryptocurrency returns using standard tools of Bayesian portfolio theory, but is not meant to be exhaustive. We discuss

[10] Alternatively, model uncertainty and robustness of this form can also be considered a model of superstition about being maximally unlucky, rather than a rational framework.

below alternatives that likely require greater modifications than we have used so far.

One alternative 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, where the first step is some binary event whereby Bitcoin "goes to zero", and the alternative is 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. The simplest is to assume zero correlation, but this 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. We have mostly explored the role of beliefs, but 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).

Finally, while our paper has an explicitly normative focus on the actions implied by various beliefs, we do not take a stand over *why* investors have the priors they do, nor whether such priors are reasonable. Much of the public debate attempts to convince others that their priors are wrong, mostly to little effect. While we seek to avoid debating priors, some readers may have heuristics to not 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 component, and we assume that everyone is entitled to their priors.

6 Conclusion

It is no exaggeration to say that the size and returns of cryptocurrency markets are an embarrassment to mainstream asset pricing. The idea of an asset class with no underlying cashflows, and no prospect of them, is deeply antithetical to many standard modelling tools. Even behavioral models tend to build off mistaken reaction to signals about cash flows. When these signals are entirely absent, and investors ought to be able to understand this, the challenges of how to model investor behavior and prices are very great.

Lacking such models, much of mainstream finance has tended towards incredulity - these assets make no sense, and so only naive or confused investors purchase them. This attitude has faced an uneasy tension, as the supposedly naive and confused investors have outperformed over long horizons equity portfolios formed on the best academic advice. Meanwhile, empirical papers on the subject have mostly skirted the big picture question of why there should be a non-zero price, and thus only obliquely addressed the question of whether investors normatively should buy any. Anecdotally, we know a fair number of professors that will confess privately to owning cryptocurrencies, but few who are willing to publicly state that anybody else ought to do the same.

Some portfolio decision has to be made. Zero is a choice like any other. Just because it is the default does not make it correct. The problem does not go away if one finds current models unsatisfactory, dislikes the assumptions of Bayesian portfolio theory, doubts the sustainability of cryptocurrency returns, or many other objections. If a student asks what their cryptocurrency allocation ought to be, what answer should one give? Positive? Negative? Zero? Not sure? We argue that it is time to tackle this tension head-on. Where possible, objections to cryptocurrency should be quantified, to see what behavior they actually justify.

The behavior that is most difficult to justify under Bayesian portfolio theory is professing a strong belief that cryptocurrency is a bubble, but refusing to take any short position. Uncertainty about cryptocurrency but having an unwillingness to take even small positions is also puzzling to a lesser extent. Quantitatively it requires either very large costs or large model uncertainty plus some additional frictions. Our results imply that many investors who were initially moderately skeptical probably ought to take some small positive weight in cryptocurrencies. That said, this is far from a blanket vindication of cryptocurrency enthusiasts, whose weights are often very high relative to the numbers we compute. Importantly, a very skeptical investor who believes cryptocurrency is a bubble, and has a short position, is acting entirely consistent with our results. However, someone who professes a strong belief that cryptocurrency is a bubble, but is unwilling to take either a short *or a long* position, at any point past or future, at any weight at all, has some explaining to do, at least if he claims to be following Bayesian portfolio theory.

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.

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 semi-groups 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.

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.

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.

Goodfellow, Ian J, Jonathon Shlens, and Christian Szegedy (2014). “Explaining and harnessing adversarial examples”. In: *arXiv preprint arXiv:1412.6572*.

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*.

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.

Iyer, Tara (2022). *Cryptic connections: spillovers between crypto and equity markets*. International Monetary Fund.

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.

Makarov, Igor and Antoinette Schoar (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 nonstockholders". 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*.

Samuelson, Paul (1963). "Risk and uncertainty: A fallacy of large numbers". In: *scientia* 57.98.

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

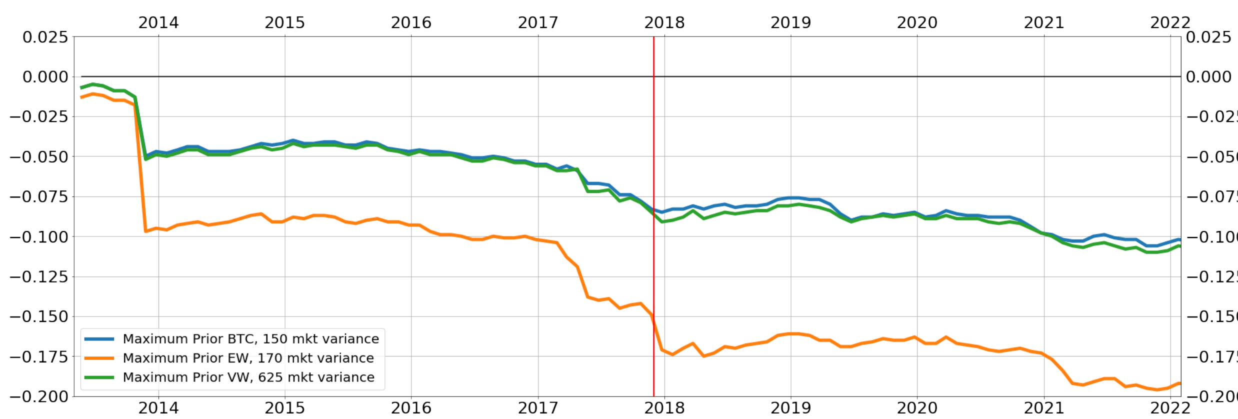

Panel B: Cumulative Non-Investment

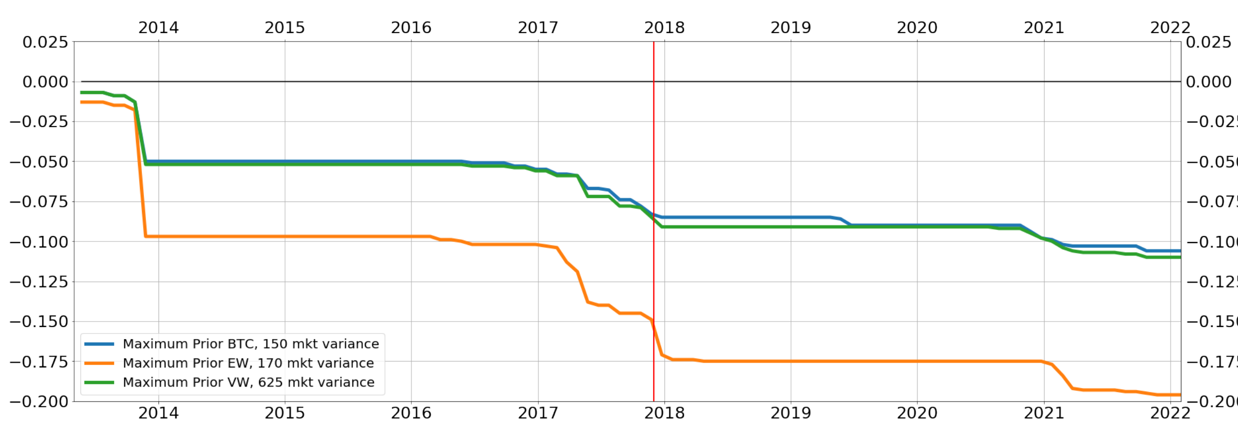

Figure 2: Optimal Cryptocurrency Portfolio Weights

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 forBitcoin, 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

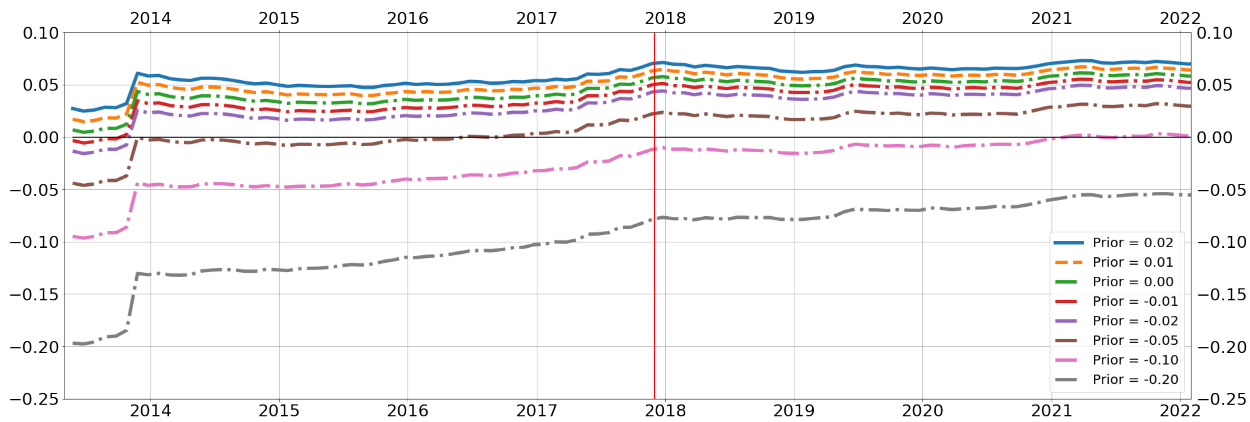

Panel B: Equal-weighted Cryptocurrency Portfolio

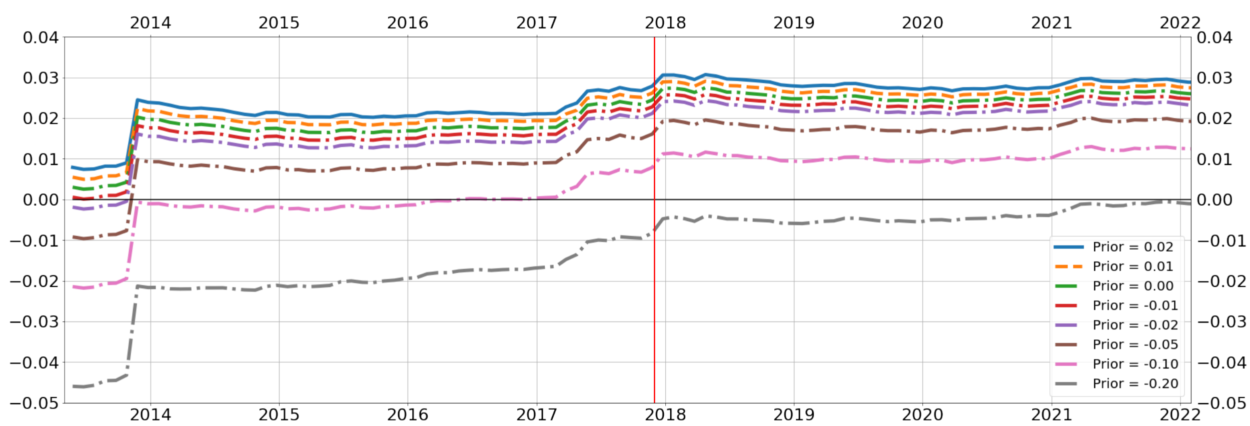

Panel C: Maximum absolute weights of BTC

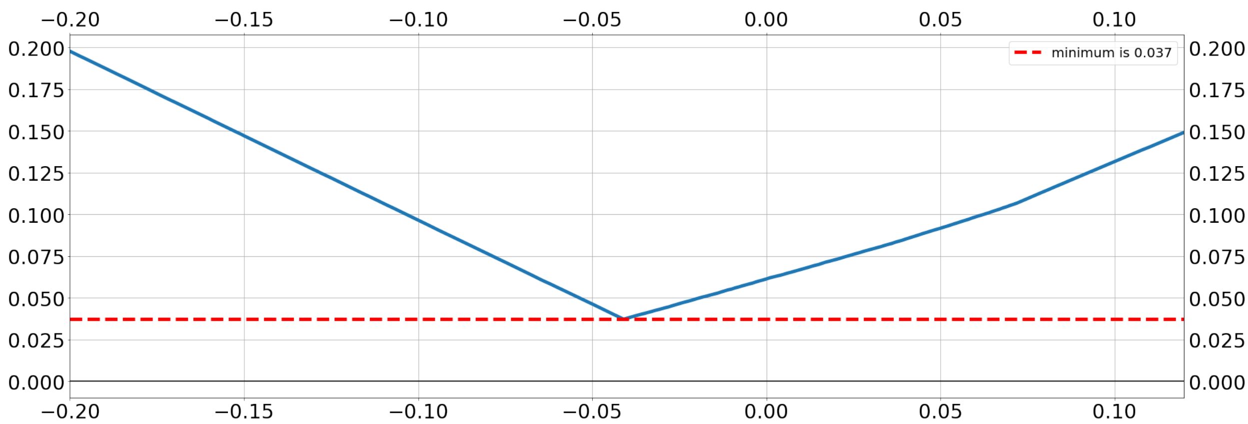

Panel D: Maximum absolute weights of Equal-weighted Cryptocurrency Portfolio

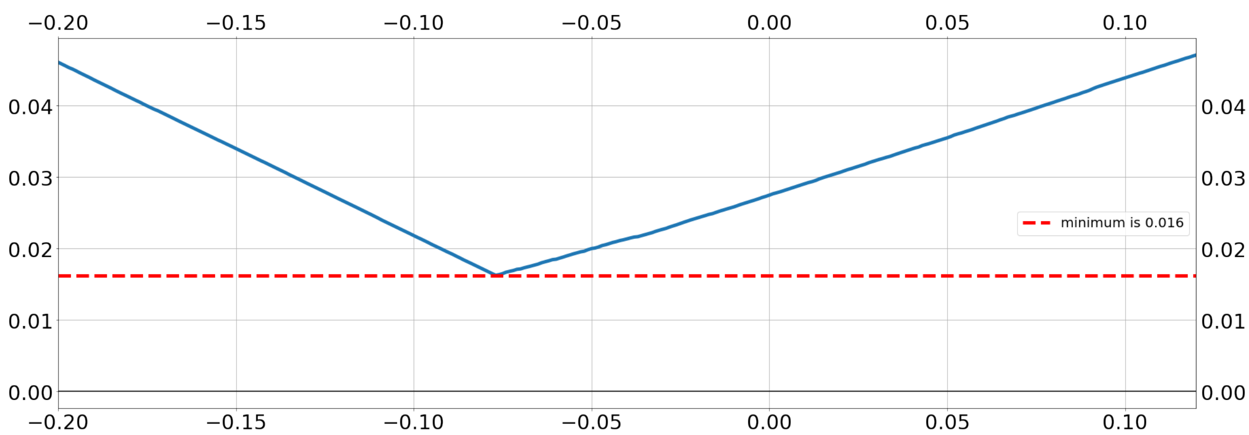

Panel E: Maximum absolute weights of BTC, no short before 2017

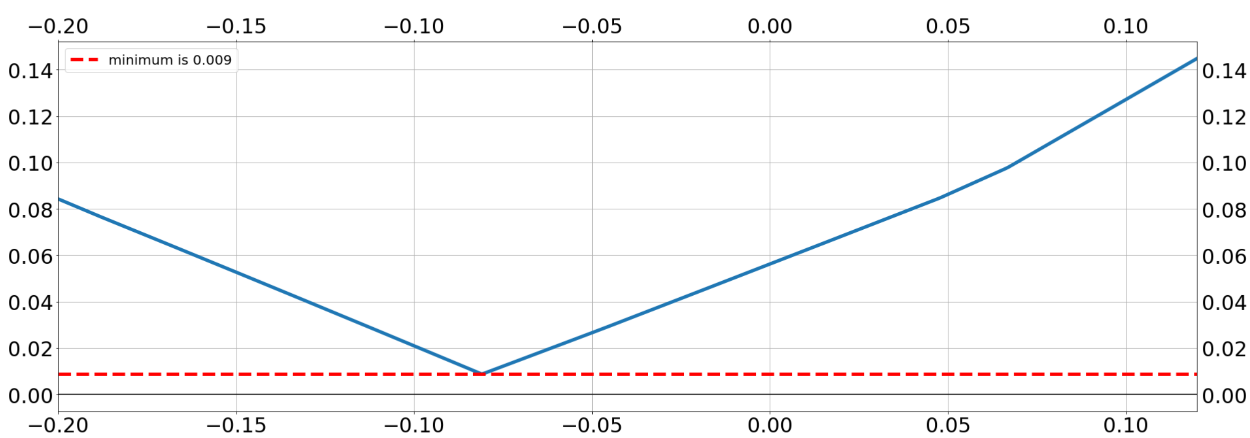

Panel F: Maximum absolute weights of Equal-weighted Cryptocurrency, no short

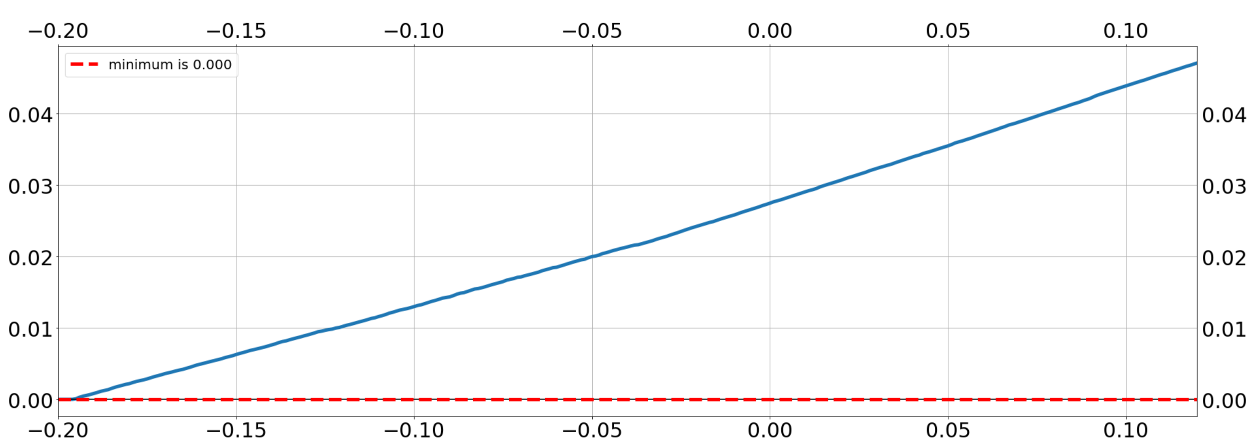

Figure 3: Ex-Ante Certainty Equivalent Gains from Cryptocurrency

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

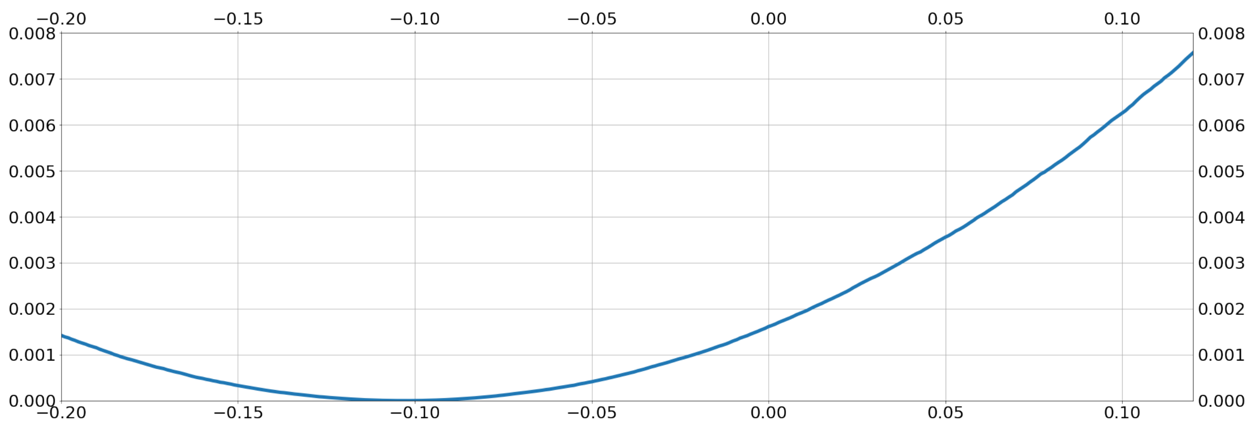

Panel B: Equal-weighted Cryptocurrency Portfolio

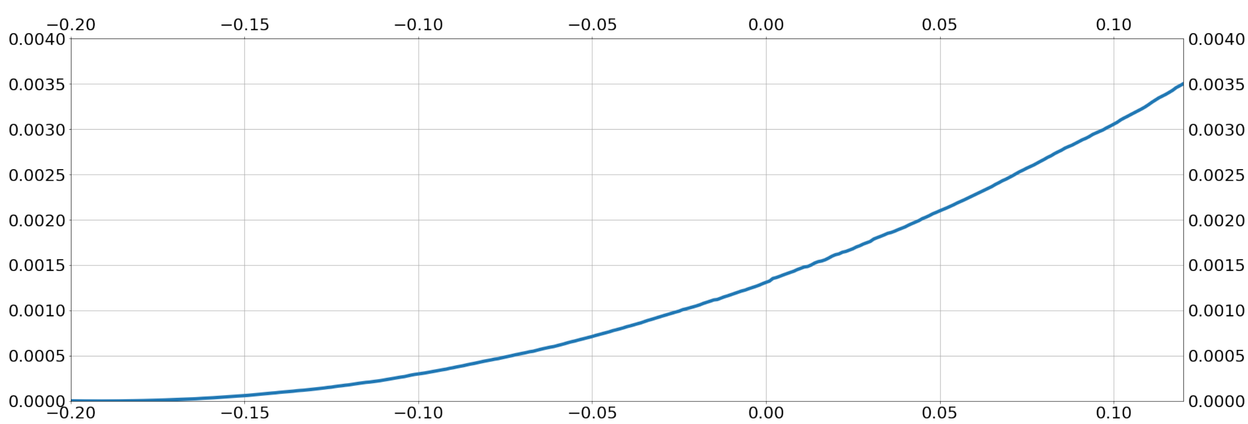

Figure 4: Ex-Post Certainty Equivalent Gains from Cryptocurrency

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

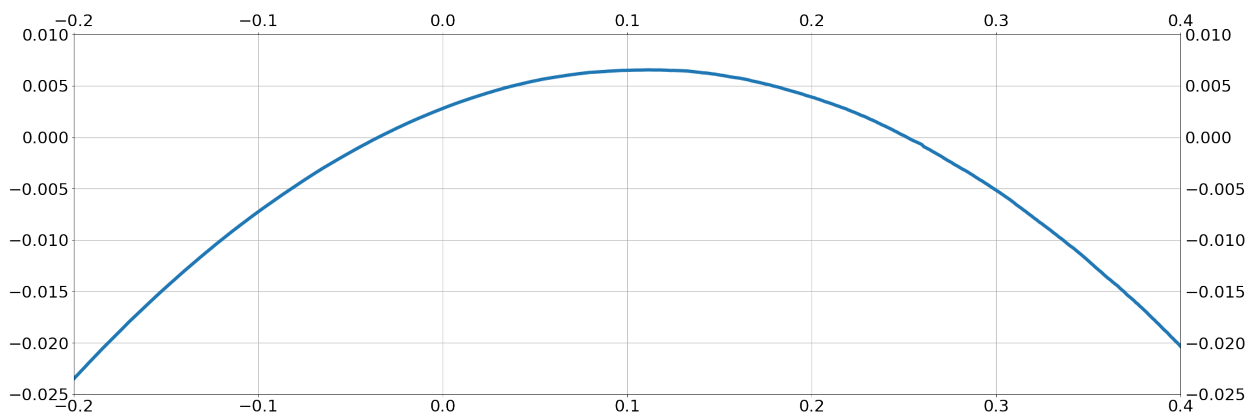

Panel B: Equal-weighted Cryptocurrency Portfolio

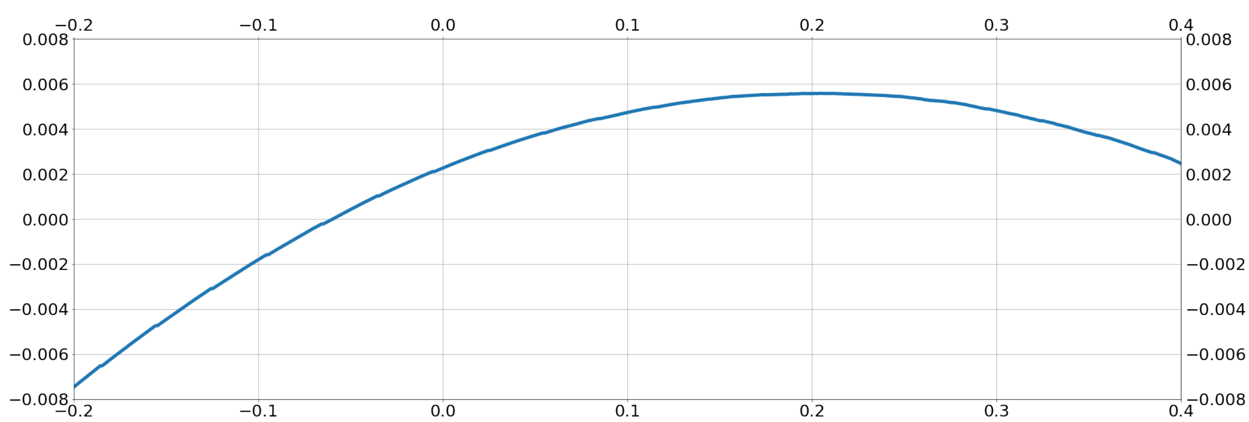

Figure 5: Ex-ante vs. Ex-post Performance of Cryptocurrency

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

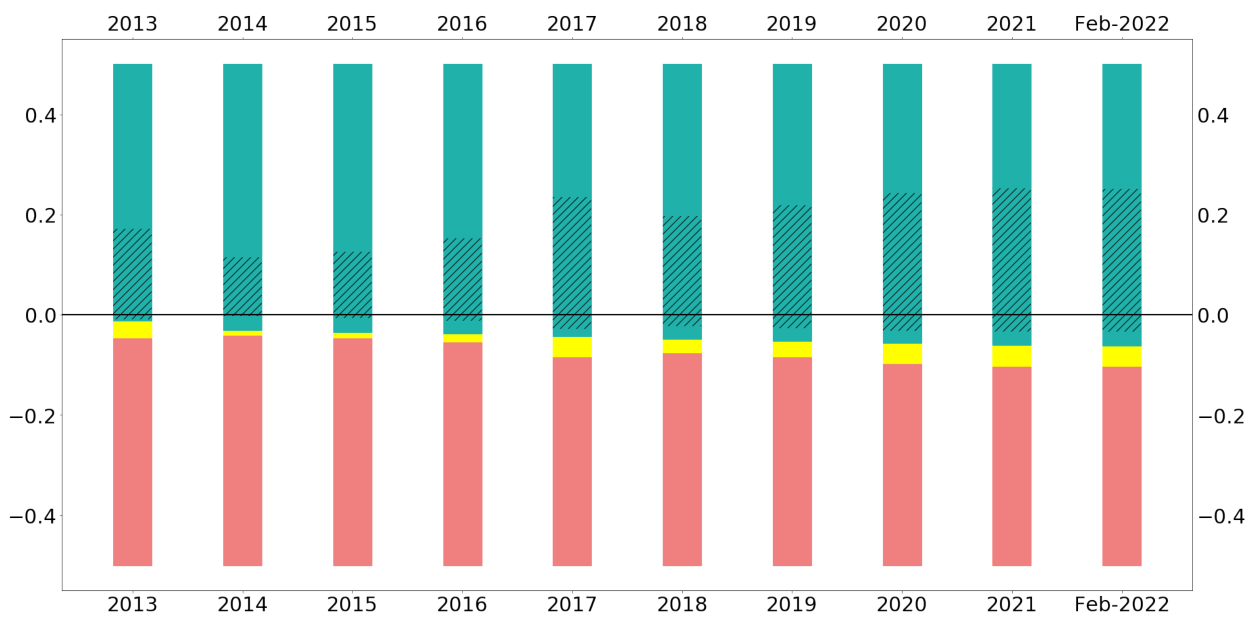

Panel B: Equal-weighted Cryptocurrency Portfolio

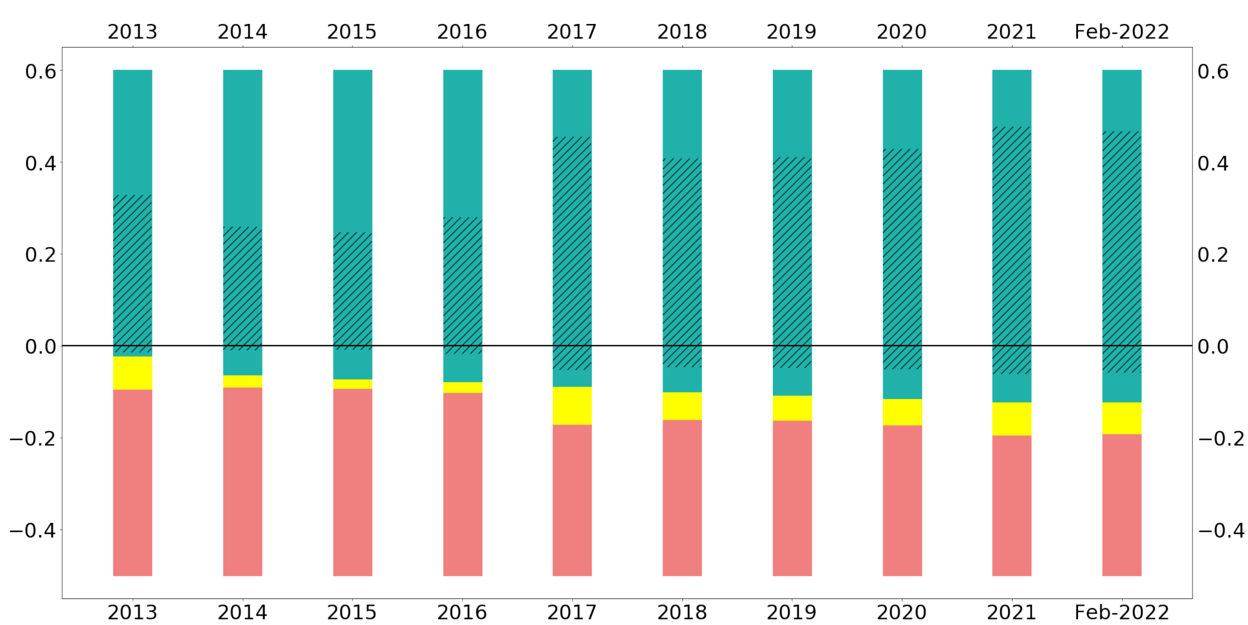

Table 1: Summary Statistics

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	0.081	0.056	0.060	0.069	0.104	0.084	0.093	0.105	0.111	0.112
Positive CER, max prior T0=10	0.171	0.115	0.125	0.152	0.235	0.197	0.219	0.243	0.253	0.251
Positive CER, min prior T0=10	-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
Positive CER, max prior T0=10	0.328	0.259	0.247	0.279	0.455	0.408	0.410	0.429	0.477	0.467
Positive CER, min prior T0=10	-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: Final 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

7.1 Towards a Theory of Bitcoin Prices

In this paper, we have been agnostic about the basis that people may have for their prior beliefs about cryptocurrency returns. In particular, 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 the Bitcoin they hold will somehow produce a dividend denominated in US dollars. Whatever is driving price changes is clearly not the cash flows, which are zero, but something else. So what might that something else be?

Our arguments here are not original in the Bitcoin community, but it may be useful to summarize for an asset pricing audience some of the more credible cases for non-zero Bitcoin prices. Necessarily, these are arguments that are loose, polemic, and proceed by analogy rather than as formal equilibrium models. We offer them as starting points for further thinking, especially for readers who may be skeptical of our findings without some kind of argument, even if not fully fleshed out, of what the alternative economic basis for the prices may be. Even if these arguments are wrong (which they may be) or incomplete (which they are almost by definition), we do not think this detracts from the main point of the paper. Something very large is missing in our models of Bitcoin prices, and it seems appropriate to have some agnosticism over what that something is.

Broadly speaking, 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.

Of these, 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. Indeed, what is striking is that a great many of the criticisms that skeptics level 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 and 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. These bars then proceed to largely just sit there.[11]. There is a huge amount of metal in those vaults. Many of the bars have sat 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? The gold produces no cash flows, but can only be sold to a new buyer.

This suggests various questions. Is gold a bubble? It depends exactly what one means by this. There is surely some industrial and jewelry demand, so the equilibrium price is unlikely to be zero. The challenge is whether these factors are the only or the main driver of gold prices. If the question about gold being a bubble means "is the price of gold substantially higher than it would be if all the central banks decided one day 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 at random 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 underlying difference between gold and Bitcoin is that gold has fundamental sources of demand from industrial uses and jewelry. In this sense, a tomato also does not produce underlying cash flows, except it is something that consumers demand. Gold, unlike Bitcoin, has potential end users of the product. Industrial uses are the most straightforward, but these also cover a very small fraction of gold's history, and so surely cannot be the underlying explanation. Rather, the fundamental explanation that makes

[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 https://www.newyorkfed.org/aboutthefed/goldvault.html for NY Fed numbers, https://www.gold.org/goldhub/data/how-much-gold for gold reserves, 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 bars underground have any sense at all, is the wedding ring. Gold is valuable, so the theory goes, because people want to put it in jewelry.

From this point of view, one of the underappreciated aspects of genius in Nakamoto 2008 (on top of the idea of the blockchain, distributed trustless consensus and solving the double-spend problem) was the realization that the wedding ring may actually be a relatively unimportant aspect of gold. Indeed, it is also worth pondering the extent to which economists can rule out the possibility that causality is actually reversed - people desire gold for wedding rings because it is valuable, not the other way around.

Rather, the most important aspect of gold may instead be that it is in fixed supply. This means that when the price increases, more cannot be easily mined to push the price back down. If some investors have extrapolative beliefs (such as diagnostic expectations in Bordalo, Gennaioli, and Shleifer 2018), then these price increases lead to more investor demand. There assuredly has to be some initial source of demand to get the process going. It may come from initial true believers in the project, for instance, or believers in Austrian economics and hard money.

It also seems useful that there is some narrative that makes the demand understandable - wedding rings form a very convenient and reasonable explanation for gold, even if the vast majority of trading comes from central banks and speculators. For Bitcoin, initial narratives centered on the demand from drug dealers, online poker players, and other people wishing to avoid mainstream banks and financial reporting. But at some point, most people buying are not drug dealers, or even thinking about drug dealers - they are just buying as a bet on Bitcoin prices.

There is also a useful aspect of Bitcoin as a hedge asset. This is sometimes framed as it being an inflation hedge, or a hedge against market downturns - that is to say, claims about the properties of Bitcoin returns, as if those were a fixed fact of the universe. Rather, we think 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 to be holding when one reaches 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 24 words forming the passphrase to a private key on a hardware wallet can be stored in one's head, and records of the coins live on computers all over the world. Like gold,Bitcoin is only a hedge against intermediate disasters - catastrophe to you personally, or in one country, but with computers and the internet still operating. But 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 still elides over the biggest point - if fundamental value does not determine prices, what does? The somewhat tautological answer presented 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. In particular, 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 - they are generally around it, but may deviate when particular psychological biases or frictions lead to errors - the standard behavioral finance view H2: Prices can be anything at all, unless there is some specific force constraining them to be in a particular range.

We think H2 has a surprising amount to recommend about it. In particular, fundamental value is an example of one such force constraining prices, inasmuch as when prices deviate too far from fundamental value, it is possible to generate 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 just collect the cash flows. If one is patient and not subject to investors withdrawing capital, this strategy has considerable appeal, although it may take a long time to work. If the asset trades a long way 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.

When framed this way, it is apparent that neither of these is especially fast, and the latter has considerable implementation risks. Fundamental value can lead to other strategies, of course, but these mostly involve other traders changing their minds and deciding to pay fundamental value for the asset. This may happen, but without a theory of what the mistake is, it is hard to know when or if this should happen. These versions are both a bet on the existence of mispricing, and a bet on when it will correct itself.

Fundamental value does not create obvious strategies for Bitcoin - one can fork the code, and even the state of the ledger (as coins such as Bitcoin Cash did), but one faces an uphill battle to overcome incumbency, and in supporting a forked version one is spending valuable hash power in the meantime on a coin that will ultimately fail.

In this version, prices change because of trades and price pressure. What aspects of Bitcoin are important in this framing? The fixed supply is obviously important, as noted, 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 terminology of Bitcoiners). This aspect rather resembles central banks attitudes to gold, incidentally. However,Bitcoin's much greater uncertainty, and much more fluctuating investor base, make it much more volatile. In addition, 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 likely to 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 the other arguments, #2 and #3. Both involve a belief in Austrian economics, that either governments will realize the virtues of fixed monetary policy, or they 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 growing size and transaction rate limit of the blockchain are well known, although the advent of Layer 2 solutions (such as the lightning network for Bitcoin, and a variety of platforms for Ethereum) may circumvent this. Perhaps more importantly, it seems likely that if Bitcoin was credibly likely to replace US dollars, it would be viewed as an enormous threat to the US government, inasmuch as it threatens the ability to raise arbitrary amounts of 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 recently in light of 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. It is true for most modern assets that financial economists study, including stocks and currencies. The only difference is that for most assets, the computers are amenable to the control of foreign governments. If the question becomes "what is the largest, most liquid electronically traded asset that cannot be seized by foreign governments during a crisis?", the answer is "Bitcoin". In this sense, it seems quite possible that more central banks may at some point begin to purchase cryptocurrency in earnest - not necessarily for the same reason as El Salvador, as a means of making them an everyday currency, but just as hard asset reserves.

These arguments, if informal, nonetheless 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 (like 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 considerably greater than the current Bitcoin market capitalization of $387 billion (as of August 29th, 2022). This does not guarantee that all the same demand from the former will move to the latter, of course, 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 the oldest and largest holders of Bitcoin is likely to be important. There is no particular reason to imagine that such holders may change their mind and sell soon (since such transactions would likely reveal their identity), but should such changes occur, this would be a significantly negative sign, given the very large amounts of Bitcoin they hold, and the potential for self-fulfilling beliefs to unravel.

Fourth, new technological developments are likely to be interesting primarily 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, and perhaps most importantly, there are many ways in which the self-reinforcing cycle of belief in Bitcoin's value as an asset could also unwind. All of the above is not an argument that Bitcoin *will* succeed, merely that it *might*.