

The Journal of
Alternative
Investments

Spring 2022 – Volume 24 Issue 4 | JAI.pm-research.com

The Role of Cryptocurrencies in Investor Portfolios

Megan Czasonis, Mark Kritzman,
Baykan Pamir, and David Turkington



The Global Mark of Distinction in Alternative Investments

The Role of Cryptocurrencies in Investor Portfolios

Megan Czasonis, Mark Kritzman, Baykan Pamir,
and David Turkington

Megan Czasonis

is a managing director and head of portfolio management research at State Street Associates in Cambridge, MA.
mczasonis@statestreet.com

Mark Kritzman

is a senior lecturer at MIT Sloan School of Management and CEO of Windham Capital Management in Cambridge, MA.
kritzman@mit.edu

Baykan Pamir

is a junior trader at Alphachain Capital in London, UK.
bpamir@alum.mit.edu

David Turkington

is the head of State Street Associates in Cambridge, MA.
dturkington@statestreet.com

KEY FINDINGS

- Full-sample correlations reveal little about the diversification potential of an asset class because they do not distinguish upside from downside correlations, nor do they consider the magnitude of returns.
- Correlations estimated from shorter-interval returns do not necessarily reflect the co-movement of longer-interval returns.
- To evaluate the diversification potential of cryptocurrencies, investors should consider the direction and magnitude of returns over the short term, and the extent to which cryptocurrencies move synchronously with, or drift apart from, other asset classes over the long term.

ABSTRACT

The role of cryptocurrencies as a vehicle for speculation has been well established. However, it is less clear if cryptocurrencies can also serve to manage risk. The authors seek to determine the diversification potential of cryptocurrencies both for short and long horizons. For short horizons, they estimate correlations that consider the direction and magnitude of returns for relevant asset classes, rather than focusing on full-sample correlations, as is customary. For long horizons, they compute “single period correlations” that capture the extent to which cryptocurrencies move synchronously with, or drift apart from, other assets over an investor’s horizon. They also identify utility-maximizing allocations to cryptocurrencies directly from historical return samples that account for all features of the data as well as more nuanced preferences than are typically assumed.

Cryptocurrencies have attracted widespread interest among private investors and even some hedge funds as a vehicle for speculation. However, it is less clear what role, if any, cryptocurrencies should play in institutional portfolios. Some have argued that cryptocurrencies might serve as a hedge against crises, much the way many investors perceive gold. Others suggest that cryptocurrencies present a lottery opportunity or perhaps function as an option on unknown unknowns. Our purpose in this paper is to evaluate cryptocurrencies from a purely empirical perspective to determine their potential for diversification and to explore what type of preferences might support allocation to cryptocurrencies within institutional portfolios.

But diversification is a complex topic, even for traditional asset classes. It is not enough to calculate the correlation of daily or monthly returns, for example, and conclude that asset classes with low average correlations are good diversifiers while those with high correlations are not as good. It is important to focus on the correlation of an asset class with other asset classes when they are performing

poorly separately from when they are performing favorably, as well as the magnitude of these directional returns.

Even if an asset class displays a favorable correlation profile based on short-term returns, most investors care as much, if not more, about how asset classes interact over long horizons. An asset class that moves independently or even inversely with a portfolio's main growth engine on a daily or monthly basis, but which also drifts in the same direction over the course of the investor's investment horizon, does little to mitigate adverse long-term performance.


Unfortunately, the longer-horizon interaction of asset classes is difficult to measure for two reasons. First, if the autocorrelations of either asset class or the cross-correlations between them are non-zero at any lag, correlations based on shorter-interval returns will misestimate the correlation of longer-interval returns.¹ Second, the correlation of longer-interval returns, whether they are estimated as independent observations or overlapping observations, vary through time. We must therefore employ a new measure of co-movement called a single period correlation to capture the longer-term interaction of asset classes.

We organize the rest of the paper as follows. In the next section, we address high frequency diversification by analyzing daily and monthly returns. We first describe how we measure conditional correlations in a way that controls for the mathematical bias that occurs from eliminating a non-random part of the return sample. We then present conditional correlation matrixes that comprise four asset classes: US stocks, US Treasury bonds, gold, and cryptocurrencies. In the third section, we turn our attention to longer-horizon correlations. We begin by presenting evidence of the extent to which shorter-horizon correlations misestimate longer-horizon correlations. We then describe mathematically the relation between the correlation of shorter-interval returns and the correlation between longer-interval returns. Next, we provide evidence of the time variation of correlations estimated from longer-interval returns, thereby disqualifying them as a reliable measure of the diversification potential of an asset class. We resolve these issues by describing a new measurement of co-movement called a single period correlation, and we describe how it is related to the Pearson correlation. We then estimate this correlation for US stocks and cryptocurrencies based on monthly, annual, and triennial return intervals. Then in the fourth section, we employ a procedure called full-scale optimization, which implicitly considers correlation asymmetry as well as complex investor preferences, to illustrate the optimal allocation to cryptocurrencies across a range of expectations and preferences. We summarize the paper in the final section.

CONDITIONAL CORRELATIONS FOR INVESTORS WITH SHORT HORIZONS

Correlations, as typically measured over the full sample of returns, often misrepresent an asset class's potential to diversify other asset classes in a portfolio when diversification is most needed—that is, when the portfolio's main growth engine delivers significant losses. Moreover, upside diversification is undesirable. Investors should seek unification on the upside. Therefore, to evaluate the relevant diversification potential of cryptocurrencies, we must condition our estimate of their correlations on return samples when other important asset classes perform poorly.

¹Previous research has documented the effect of serial correlation on risk and correlation estimates. For example, Asness et al. (2001) show that simple betas and correlations estimated from monthly returns significantly understate hedge fund equity market exposure. Similarly, Anson (2002, 2007, and 2013) shows that private equity valuations exhibit considerable exposure to lagged public market returns.

EXHIBIT 1**Full Sample and Partial Sample Returns**

But doing this is not straightforward. Correlations estimated from sub-samples are biased even if the returns are jointly normally distributed, owing to the mathematics of correlations.

Consider a joint normal distribution with equal means of 0%, equal volatilities of 15%, and an unconditional (full sample) correlation of 50%. Suppose we condition the correlation only on returns below a certain threshold for the x variable, as illustrated in the right panel of Exhibit 1.

The correlation of this subset of observations is only 20%. This reduction arises because opposing extreme values are more highly correlated than observations that are more like each other. When we exclude the opposing extreme values in our calculation, the correlation goes down, but this does not suggest that the variables co-move less when the x variable's values are below the chosen threshold. To determine if an asset class offers more or less diversification than suggested by its full-sample correlation when the other asset class is performing poorly, we must account for this bias.² We do so by first simulating the change in the correlation from a full sample to a partial sample assuming normality and given the same means, standard deviations, and correlation of the empirical full sample. We then compute the empirical difference between the full-sample correlation and the partial-sample correlation. The difference in these

EXHIBIT 2**Asset Classes and Indexes^a**

Asset Classes	Indexes
US Stocks	S&P 500 Index
Treasury Bonds	Bloomberg US Long Treasury Index
Gold	S&P GSCI Gold Index
Cryptocurrencies	Bitcoin Spot Price

NOTES: ^aAll data is obtained from Datastream. Bitcoin spot prices reflect the exchange rate on Bitstamp at 5pm EST. Prior to January 2012, we backfill Bitcoin with the exchange rates from Bloomberg.

differences measures the extent to which an asset class offers more or less diversification than its full-sample correlation indicates.

We perform the same exercise for partial samples in which a chosen asset class performs above a chosen threshold, in which case we hope to find that the correlation goes down less than the reduction that would occur normally by excluding part of the return sample.

We consider four asset classes in our analysis: US stocks, Treasury bonds, gold, and cryptocurrencies.³ We base our analysis on daily and monthly returns beginning June 30, 2011, and ending December 31, 2020. We include US stocks because investors typically employ them as the main growth engine for portfolios. We include Treasury bonds because investors include them as the main diversifying asset class for portfolios. We include gold because investors often resort to gold during periods of severe financial stress.⁴ And we include cryptocurrencies to determine what role they might provide in helping investors manage risk. Exhibit 2 lists the indexes we use to represent these asset classes.

We define the return thresholds based on US stock performance, because as we mentioned earlier, most investors deploy US stocks as their portfolios' main growth engine. Our focus, therefore, is to measure the extent to which other asset classes offer diversification when US stocks perform poorly, and unification when they perform well. We focus first on daily returns.

² See Appendix A for the closed-form solution of conditional correlations.

³ We base our empirical analysis on Bitcoin returns because Bitcoin has the longest return history and because the prices of major cryptocurrencies tend to correlate very strongly with each other.

⁴ It is worth emphasizing that we include gold due to its potential value as a hedge during periods of financial distress. We do not argue that it has a positive expected return.

EXHIBIT 3**Correlation of Daily Returns with US Stocks
US Stock Returns < 0%**

	Treasury Bonds	Gold	Crypto- currencies
Full Sample	-0.42	0.02	0.11
Normal: Daily Stocks < 0%	-0.26	0.00	0.07
Empirical: Daily Stocks < 0%	-0.34	-0.02	0.15
Empirical Minus Normal	-0.07	-0.02	0.08

Exhibit 3 shows the asset class correlations when US stocks experience negative returns.

The first row of Exhibit 3 presents correlations with US stocks based on the full sample of daily returns. The second row shows how these correlations change for normally distributed returns when we exclude that part of the sample when US stock returns are positive—what we referred to earlier as correlation bias. The third row shows how the correlations change from the full sample to the partial sample given the empirical return distributions. The final row shows the incremental change in the correlations beyond the

effect of correlation bias. A negative value indicates that diversification is greater than what is indicated by the full-sample correlation when US stock returns are negative, and a positive value indicates that there is less diversification than what the full sample indicates. After controlling for the bias, we observe that when US stocks have negative returns, Treasury bonds and gold offer greater diversification benefits to US stocks.

The left panel in Exhibit 4 summarizes conditional correlations with US stocks based on daily returns with thresholds of 0% (already shown in Exhibit 3), negative one standard deviation, and negative two standard deviations for US stock returns, and for monthly returns with the same thresholds. The right panel shows the incremental changes in the correlations beyond the effect of correlation bias.

We see that the larger the loss in US stocks, the greater the co-movement between stocks and cryptocurrencies. This relationship is generally consistent across both frequencies and is magnified in daily returns.

Exhibit 5 shows the same information as Exhibit 4, but this time for upside US stock returns. In Exhibit 5, we would like to see positive differences between the empirical values and values associated with a normal distribution.

On the upside, and especially for larger moves in daily returns, cryptocurrencies have the most favorable correlation to stocks.

**SINGLE PERIOD CORRELATIONS FOR INVESTORS
WITH LONG HORIZONS**


As we mentioned previously, investors care as much, if not more, about how the main growth and defensive components of their portfolios co-move over the full

EXHIBIT 4**Correlations with US Stocks for Downside US Stock Returns**

	Empirical			Empirical Minus Normal		
	Treasury Bonds	Gold	Crypto- currencies	Treasury Bonds	Gold	Crypto- currencies
Daily Stocks, Full Sample	-0.42	0.02	0.11	0.00	0.00	0.00
Daily Stocks < 0%	-0.34	-0.02	0.15	-0.07	-0.02	0.08
Daily Stocks < -1σ	-0.23	0.05	0.30	-0.03	0.04	0.23
Daily Stocks < -2σ	-0.15	0.09	0.37	0.00	0.11	0.28
Monthly Stocks, Full Sample	-0.41	0.07	0.18	0.00	0.00	0.00
Monthly Stocks < 0%	-0.32	-0.02	0.18	-0.08	-0.04	0.08
Monthly Stocks < -1σ	-0.27	0.17	0.22	-0.06	0.14	0.13
Monthly Stocks < -2σ	-0.25	-0.43	0.19	-0.10	-0.47	0.09

EXHIBIT 5**Correlations with US Stocks for Upside US Stock Returns**

	Empirical			Empirical Minus Normal		
	Treasury Bonds	Gold	Cryptocurrencies	Treasury Bonds	Gold	Cryptocurrencies
Daily Stocks, Full Sample	-0.42	0.02	0.11	0.00	0.00	0.00
Daily Stocks > 0%	-0.29	0.05	0.07	-0.02	0.04	0.01
Daily Stocks > 1 σ	-0.26	0.15	0.26	-0.07	0.14	0.22
Daily Stocks > 2 σ	-0.17	0.18	0.33	-0.03	0.19	0.29
Monthly Stocks, Full Sample	-0.41	0.07	0.18	0.00	0.00	0.00
Monthly Stocks > 0%	-0.14	0.26	0.02	0.14	0.22	-0.10
Monthly Stocks > 1 σ	0.27	0.14	0.03	0.47	0.10	-0.06
Monthly Stocks > 2 σ	0.61	0.48	0.40	0.74	0.42	0.35

EXHIBIT 6**Correlation of Cryptocurrencies and US Stocks (June 2011 through December 2020)**

extent of their investment horizon as they do about how they co-move daily or monthly. To address this issue, we focus on the correlation of one- and three-year returns.

The conventional approach for forecasting the correlation of one- or three-year returns is to extrapolate the correlation of monthly returns from a recent historical sample. This approach assumes that correlations are invariant to the return interval used to estimate them. However, this invariance would only be true if the returns of the asset classes were each serially independent at all lags and if their lagged cross correlations were also zero at all lags. For many asset classes these implicit assumptions are not borne out by historical evidence. Exhibit 6 offers evidence of the divergence of correlations estimated from high frequency and low frequency observations.

This divergence in correlations across different time intervals could arise if short-term returns respond similarly to a given factor, whereas long-term returns drift apart in response to a different lower frequency influence. Equation 1 shows how the correlation of longer-interval returns is mathematically related to the correlation of shorter-interval returns.

$$\rho(x_t + \dots + x_{t+q-1}, y_t + \dots + y_{t+q-1}) = \frac{q\rho_{x_t, y_t} + \sum_{k=1}^{q-1} (q-k)(\rho_{x_{t+k}, y_t} + \rho_{x_t, y_{t+k}})}{\sqrt{q + 2\sum_{k=1}^{q-1} (q-k)\rho_{x_t, x_{t+k}}} \sqrt{q + 2\sum_{k=1}^{q-1} (q-k)\rho_{y_t, y_{t+k}}} \quad (1)$$

In Equation 1, $\rho(x_t + \dots + x_{t+q-1}, y_t + \dots + y_{t+q-1})$ is the correlation of the cumulative continuous returns of x and y over q periods.

The numerator of Equation 1 equals the covariance of the assets taking lagged cross correlations into account, whereas the denominator equals the product of the assets' standard deviations taking their autocorrelations into account. To apply this equation, we would need to estimate the autocorrelations of stocks and bonds at all lags as well as their lagged cross correlations at all lags, which would be unduly cumbersome.⁵

Alternatively, we can calculate the correlation of independent one- or three-year returns, but this approach has two drawbacks. First, the estimate of the correlation is highly sensitive to the start date of the first observation. And second, the correlations are unlikely to be constant across a long history.

These problems might lead us to consider using overlapping observations to calculate longer-horizon correlations—but this approach, while mitigating the start-date problem, is still subject to issues of time variation.

Therefore, we need an approximation of the correlation of returns that captures the extent to which these returns move synchronously or drift apart during each one- or three-year period. Just as each single return observation for an asset class differs from its long-run average returns, the co-movement of asset classes during each one- or three-year period differs from how they co-move on average over the long run.

Equation 2 captures the co-movement of the cumulative return of two asset classes for a chosen period considering the average co-movement of the asset classes over the full sample of those periodic returns.⁶

$$\text{single period correlation}_t(x, y) = \frac{\left(\frac{r_{x,t} - \mu_x}{\sigma_x}\right) \left(\frac{r_{y,t} - \mu_y}{\sigma_y}\right)}{\frac{1}{2} \left(\left(\frac{r_{x,t} - \mu_x}{\sigma_x}\right)^2 + \left(\frac{r_{y,t} - \mu_y}{\sigma_y}\right)^2 \right)} \quad (2)$$

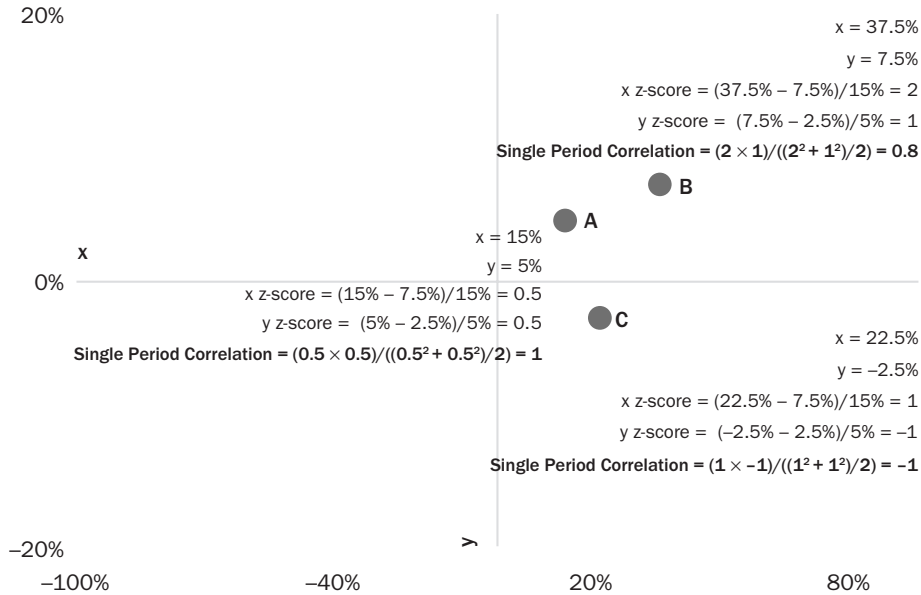
In Equation 2, r_x and r_y equal the cumulative return of the two asset classes over a chosen period, μ_x and μ_y equal the long-run arithmetic average return of the asset classes, and σ_x and σ_y equal the standard deviation of returns with the same periodicity.

We refer to this measure as a single period correlation, because it measures the co-movement of the cumulative returns of the asset classes for a single period rather than the average correlation of those periodic returns over the full sample. Each single period correlation will differ from the long-run average correlation based on the period specific co-movement of the asset classes.

For example, as shown in Exhibit 7, if the z-scores of the two asset classes are identical for the period, the single period correlation will equal 1. If they are exactly opposite, the single period correlation will equal -1 . If they have the same sign but different magnitudes, the single period correlation will be between 0 and 1. And if they have different signs and different magnitudes, the single period correlation will be between -1 and 0. It is important to note that even though this measure relies on means and standard deviations estimated from the full sample of all returns, it informs us about the co-movement of the two asset classes for a single period.

EXHIBIT 7

Illustration of Single Period Correlation



Average return of x and y = 7.5% and 2.5%, respectively; Standard deviation of x and y = 15% and 5%, respectively.

The single period correlation is related to the correlation of a time series of asset class returns, which is called the Pearson correlation, in a precise mathematical way. It equals a weighted average of the single period correlations throughout the sample, where the weights equal the informativeness of each period's returns.

$$Correlation(x, y) = \frac{1}{N} \sum_{t=1}^N (informativeness_t \times single\ period\ correlation_t) \quad (3)$$


Here, N equals the number of return periods in the sample and $informativeness_t$ and $single\ period\ correlation_t$ are based on asset class returns ending at time t . In this context, informativeness is defined as the average of the squared z-scores of asset class returns for a given period. This calculation acknowledges that patterns of co-movement among variables are more meaningful when their magnitudes are large, as those observations are more likely to reflect events as opposed to noise.

If, in Equation 3, we re-write $informativeness$ as the average squared z-score of asset class returns and substitute Equation 2 for $single\ period\ correlation$, the correlation then becomes:

$$Correlation(x, y) = \frac{1}{N} \sum_{t=1}^N \left(\frac{1}{2} \left(\left(\frac{r_{x,t} - \mu_x}{\sigma_x} \right)^2 + \left(\frac{r_{y,t} - \mu_y}{\sigma_y} \right)^2 \right) \times \frac{\left(\frac{r_{x,t} - \mu_x}{\sigma_x} \right) \left(\frac{r_{y,t} - \mu_y}{\sigma_y} \right)}{\frac{1}{2} \left(\left(\frac{r_{x,t} - \mu_x}{\sigma_x} \right)^2 + \left(\frac{r_{y,t} - \mu_y}{\sigma_y} \right)^2 \right)} \right) \quad (4)$$

This expression simplifies to the average product of asset class z-scores over the full sample, which is equivalent to the full-sample covariance of asset class returns

EXHIBIT 8
Single Period Correlations between US Stocks and Cryptocurrencies



divided by the product of their full-sample standard deviations, the formula commonly used to measure a time series correlation.

$$Correlation(x, y) = \frac{1}{N} \sum_{t=1}^N \left(\left(\frac{r_{x,t} - \mu_x}{\sigma_x} \right) \left(\frac{r_{y,t} - \mu_y}{\sigma_y} \right) \right) = \frac{Cov(x, y)}{\sigma_x \sigma_y} \quad (5)$$

The bottom line is that the single period correlation is a better representation of the correlation of one- or three-year returns than the correlation of monthly returns because it captures the drift of asset class returns whereas the correlation of monthly returns does not. And it is better than the correlation of either independent blocks of one- or three-year returns or overlapping returns because it captures the time varying nature of the correlation whereas the full-sample correlation does not.

Exhibit 8 presents single period correlations between US stocks and cryptocurrencies for monthly, annual, and triennial return periods. For each return interval, we report the single period correlation for each observed return of that periodicity from June 2011 through December 2020. All observations are monthly. For reference, the dotted lines

show the Pearson correlations (also reported in Exhibit 6) which, as described previously, represent average single period correlations across the full sample.

Exhibit 8 highlights the time-varying nature of correlations. Though the average correlation between US stocks and cryptocurrencies is positive across all three return intervals, there are individual periods with highly divergent returns between the two assets. Moreover, the relationship between their three-year returns experienced a structural shift in 2018. Prior to 2018, their long-term co-movement was consistently positive; since then, it often has been negative, including over the most recent three-year period ended December 2020. Though it is difficult to determine the exact reasons for this shift, it does align with a dramatic and sustained jump in Bitcoin trading volumes. At the beginning of 2018, average Bitcoin volumes were approximately 100 times greater than at the beginning of 2017, and they have remained elevated since.⁷

FULL-SCALE OPTIMIZATION FOR INVESTORS WITH SHORT HORIZONS

In Part 2, we presented conditional correlations to indicate the potential of cryptocurrencies to diversify against adverse US stock returns and to unify with US stocks when they perform favorably. Although this information is helpful for understanding the potential of cryptocurrencies to improve a portfolio's return and risk profile, it does not explicitly reveal how much to allocate to them. We therefore employ a technique called full-scale optimization to determine the optimal allocation to cryptocurrencies given a set of assumed investor beliefs and preferences.

Unlike mean-variance analysis, which assumes that returns are either elliptically distributed or that investor preferences can be described by just mean and variance, full-scale optimization accounts for all features of the data, including how correlations vary across sub-samples, and it accounts for preferences that depend on higher moments. Accounting for these features is particularly important when modeling cryptocurrencies, given their historical propensity for extreme outliers. For example, from July 2011 through December 2020, Bitcoin's average monthly return was 13%, its volatility was 51%, and its skewness and kurtosis were 6.0 and 50.0, respectively.⁸ We would be remiss to ignore these higher moments when determining an optimal allocation to cryptocurrencies.

Here is how full-scale optimization works. We specify a utility function and a historical return sample for the asset classes we wish to consider. Next, we select a potential allocation and compute the utility of that allocation for all the periods in our return sample, given a specified utility function. We then compute the average utility across all periods. Next, we select another set of allocations and compute its average utility across all periods. We proceed in this fashion, systematically testing alternative portfolio combinations (typically about 275,000) to be reasonably confident that one of them offers the highest possible utility given the return history and utility function.

In our analysis, we model investor preferences to include both aversion to losses and a preference for speculative gains. These elements capture the two reasons an investor may choose to invest in cryptocurrencies: diversification and speculation. To model risk aversion, we assume kinked power utility, where the kink corresponds to a dramatic decline in utility for losses below a specified level. To reflect a lottery preference, we assume that utility jumps at a specified return level. We introduce a jump in preferences to reflect the notion that an investor derives new benefits beyond a certain level of wealth that would not occur incrementally.

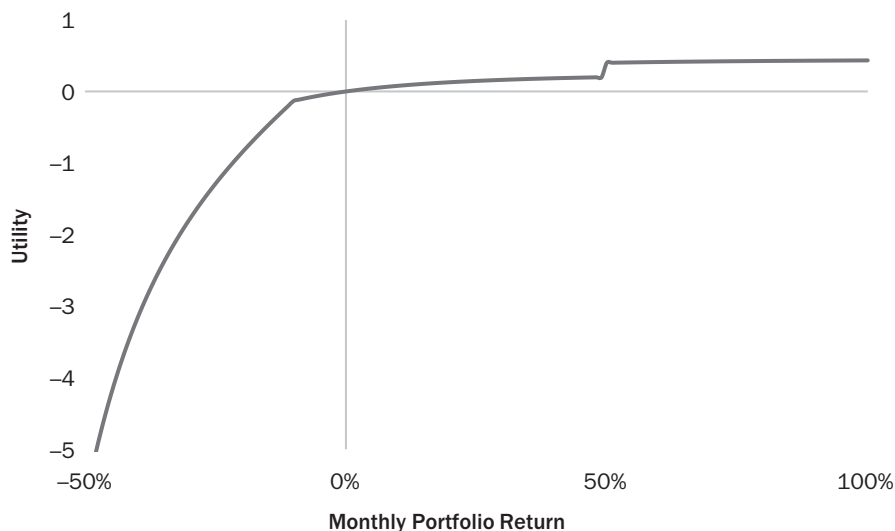
Equation 6 describes these preferences as a function of portfolio returns (x):

⁷ See Appendix B for a chart of Bitcoin trading volumes.

⁸ See Appendix C for a histogram of monthly Bitcoin returns.

EXHIBIT 9

Kinked Utility Function with Lottery Preference



$$U(x) = \begin{cases} v(x - \theta_1) + \frac{(1+x)^{1-\gamma} - 1}{1-\gamma}, & x < \theta_1 \\ \frac{(1+x)^{1-\gamma} - 1}{1-\gamma}, & \theta_1 \leq x < \theta_2 \\ j + \frac{(1+x)^{1-\gamma} - 1}{1-\gamma}, & x \geq \theta_2 \end{cases} \quad (6)$$

Here, γ is the risk aversion coefficient; θ_1 indicates the location of the risk aversion kink, v the steepness of the loss aversion slope, j the lottery jump, and θ_2 the location of the lottery kink.

Exhibit 9 shows the utility associated with different levels of portfolio return. For this illustration, we set the risk aversion coefficient (γ) equal to 5, the risk aversion kink (θ_1) equal to -10% , the risk aversion slope (v) equal to 5, the lottery jump (j) equal to 0.20, and the lottery kink (θ_2) to 50%. We calibrate these parameters to deliver roughly a 60/40 allocation to US stocks and Treasuries in the absence of cryptocurrencies.

Next, we apply full-scale optimization and the specified utility function to monthly returns for US stocks, Treasury bonds, gold, and cryptocurrencies for the period June 2011 through December 2020. We vary the degree of lottery preference by modeling a range of lottery kinks (from 20% to 90%). Moreover, we consider a range of expected returns for cryptocurrencies (from 0% to 100% per year).⁹ Exhibit 10 reports the optimal allocations to cryptocurrencies for various combinations of expected return and lottery preference. We also report the results for an investor with kinked utility and no lottery preference.

Exhibit 10 reveals that a lottery preference for returns above 20% supports some exposure to cryptocurrencies even if they are not expected to generate a positive return. Without a lottery preference, however, investors require that cryptocurrencies return at least 30% to support an allocation to them. In general and as expected, the optimal

⁹We assume expected returns of 7.5%, 4%, and 0% per year for US stocks, Treasury bonds, and gold, respectively.

EXHIBIT 10**Optimal Allocations to Cryptocurrencies**

		Location of Lottery Preference								None
		20%	30%	40%	50%	60%	70%	80%	90%	
Bitcoin Expected Return (per year)	0%	5%	0%	0%	0%	0%	0%	0%	0%	0%
	10%	5%	0%	0%	0%	0%	0%	0%	0%	0%
	20%	5%	0%	0%	0%	0%	0%	0%	0%	0%
	30%	5%	7%	1%	1%	1%	1%	1%	1%	1%
	40%	5%	7%	9%	2%	2%	2%	2%	2%	2%
	50%	5%	7%	9%	3%	3%	3%	3%	3%	3%
	60%	12%	7%	9%	11%	4%	4%	4%	4%	4%
	70%	12%	7%	9%	11%	14%	6%	6%	6%	6%
	80%	12%	7%	9%	11%	14%	16%	7%	7%	7%
	90%	12%	18%	9%	12%	14%	16%	18%	9%	9%
	100%	12%	18%	11%	12%	14%	16%	18%	20%	11%

allocation to cryptocurrencies falls with increases in the lottery threshold and rises with increases in expected return. Of course, these results are specific to our specification of utility as well as the return sample we used for our analysis.

SUMMARY

We evaluated the role of cryptocurrencies in investor portfolios. We first considered the diversification potential of cryptocurrencies for investors with short horizons. We argued that an unconditional correlation might not give a good indication of an asset's diversification potential because it fails to consider the direction and magnitude of returns. We therefore reported correlations conditioned on the performance of US stocks, because investors usually include US stocks as the main growth engine of their portfolios. Although cryptocurrencies performed well when US stocks performed well, they also moved in tandem with US stocks when stocks performed poorly. This correlation profile suggests that cryptocurrencies do not offer protection for investors with short horizons.

We then considered the diversification potential of cryptocurrencies for investors with long horizons. We offered evidence that correlations estimated from monthly returns misrepresent the correlation of annual and three-year returns. And we argued that average correlations estimated from longer-interval returns are unreliable because they vary through time. We therefore introduced the single period correlation to measure the extent to which returns move synchronously or drift apart over a particular long horizon. Based on this analysis, we showed that, although cryptocurrencies drifted in the same direction as US stocks over three-year intervals up to 2017, since then their co-movement has fluctuated between drifting together and drifting apart. During the most recent three-year period in our sample (January 2018–December 2020), cryptocurrencies and US stocks drifted apart. Though we do not know the exact reasons for this shift between 2017 and 2018, it does align with a dramatic and sustained jump in Bitcoin trading volumes.

Finally, we employed a portfolio formation technique called full-scale optimization to determine the optimal allocation to cryptocurrencies, given historical returns, various expected returns, and a utility function that assumes investors are especially averse to losses below a specified threshold but have lottery preferences above other thresholds. This approach to portfolio construction considers every feature of the return sample as well as the nuances of a complex utility function. Our analysis

revealed that a lottery preference was sufficient to support some allocation to cryptocurrencies as long as their expected return was not negative. However, without a lottery preference, investors require an annualized expected return of at least 30% to support allocation to cryptocurrencies.

Our analysis is based on a single pass through history; therefore, one should not draw immutable conclusions from it. Instead, we suggest that investors treat our analysis as a template for continued empirical evaluation of this important topic.

Also, it is important to keep in mind that the inferences we draw are based purely on statistical patterns from the perspective of a single investor. We did not consider how cryptocurrencies might function within a broader equilibrium framework, which is the perspective from which policymakers should consider the role of cryptocurrencies. But such a perspective is beyond the scope of our analysis.

APPENDIX A

CLOSED-FORM SOLUTION OF CONDITIONAL CORRELATION FOR BIVARIATE NORMAL DISTRIBUTION¹⁰

Let $X = (x, y) \sim N(0, \Sigma)$ where Σ has unit variances and unconditional correlation ρ . We define the correlation of x and y conditional on observations for which $x < h$ as

$$\text{Corr}(x, y|x < h) = \frac{\text{Cov}(x, y|x < h)}{\sqrt{\text{Var}(x|x < h) \text{Var}(y|x < h)}} = \frac{m_{11} - m_{10}m_{01}}{\sqrt{(m_{20} - m_{10}^2)(m_{02} - m_{01}^2)}}$$

Let $L(\cdot)$ denote the cumulative density of a truncated bivariate normal distribution, then

$$\begin{aligned} L(h) &= \int_{-\infty}^h \int_{-\infty}^{\infty} \frac{1}{2\pi \sqrt{1-\rho^2}} e^{-\frac{x^2-2\rho xy+y^2}{2(1-\rho^2)}} dx dy \\ &= \int_{-\infty}^h \int_{-\infty}^{\infty} \frac{1}{2\pi \sqrt{1-\rho^2}} e^{-\frac{x^2-2\rho xy+y^2+\rho^2 x^2-\rho^2 x^2}{2(1-\rho^2)}} dx dy \\ &= \int_{-\infty}^h \int_{-\infty}^{\infty} \frac{1}{2\pi \sqrt{1-\rho^2}} e^{-\frac{(y-\rho x)^2+(1-\rho^2)x^2}{2(1-\rho^2)}} dx dy \\ &= \int_{-\infty}^h \int_{-\infty}^{\infty} \frac{1}{2\pi \sqrt{1-\rho^2}} e^{-\frac{(y-\rho x)^2}{2(1-\rho^2)} - \frac{x^2}{2}} dx dy \\ &= \int_{-\infty}^h \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi} \sqrt{1-\rho^2}} e^{-\frac{(y-\rho x)^2}{2(1-\rho^2)}} dy \end{aligned}$$

With a change of variables, $z = \frac{y - \rho x}{\sqrt{1 - \rho^2}}$ and $dz = \frac{dy}{\sqrt{1 - \rho^2}}$, the equation above simplifies to

$$\begin{aligned} L(h) &= \int_{-\infty}^h \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{z^2}{2}} dz \\ &= \int_{-\infty}^h \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx \end{aligned}$$

¹⁰See Ang et al. (2002) for a closed-form solution for estimating conditional correlation for a more general case when both assets' returns fall within a specified range.

Going forward, we will denote $L(h)$ as $\Phi(h)$. Next, we derive the moments. By definition, m_{i0} are:

$$m_{i0} \Phi(h) = \int_{-\infty}^h \frac{x^i}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx$$

For $i = 1, 2$:

$$\Phi(h) m_{10} = -\frac{1}{\sqrt{2\pi}} e^{-\frac{h^2}{2}}$$

$$\Phi(h) m_{20} = \Phi(h) - \frac{h}{\sqrt{2\pi}} e^{-\frac{h^2}{2}}$$

$$\begin{aligned} \Phi(h) m_{01} &= \int_{-\infty}^h \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx \int_{-\infty}^{\infty} \frac{y}{\sqrt{2\pi} \sqrt{1-\rho^2}} e^{-\frac{(y-\rho x)^2}{2(1-\rho^2)}} dy \\ &= \int_{-\infty}^h \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} (\rho x) dx = -\frac{\rho}{\sqrt{2\pi}} e^{-\frac{h^2}{2}} \end{aligned}$$

$$\begin{aligned} \Phi(h) m_{02} &= \int_{-\infty}^h \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx \int_{-\infty}^{\infty} \frac{y^2}{\sqrt{2\pi} \sqrt{1-\rho^2}} e^{-\frac{(y-\rho x)^2}{2(1-\rho^2)}} dy \\ &= \int_{-\infty}^h \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} (1-\rho^2 + \rho^2 x^2) dx \\ &= (1-\rho^2) \Phi(h) + \rho^2 \int_{-\infty}^h \frac{x^2}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx \\ &= (1-\rho^2) \Phi(h) + \rho^2 \left[\Phi(h) - \frac{h}{\sqrt{2\pi}} e^{-\frac{h^2}{2}} \right] \\ &= \Phi(h) - \frac{\rho^2 h}{\sqrt{2\pi}} e^{-\frac{h^2}{2}} \end{aligned}$$

$$\begin{aligned} \Phi(h) m_{11} &= \int_{-\infty}^h \frac{x}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx \int_{-\infty}^{\infty} \frac{y}{\sqrt{2\pi} \sqrt{1-\rho^2}} e^{-\frac{(y-\rho x)^2}{2(1-\rho^2)}} dy \\ &= \int_{-\infty}^h \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} \rho x^2 dx = \rho \left[\Phi(h) - \frac{h}{\sqrt{2\pi}} e^{-\frac{h^2}{2}} \right] \end{aligned}$$

In this final step we substitute the moments in and obtain the conditional variances and covariance.

$$\text{Var}(x|x < h) = m_{20} - m_{10}^2 = \frac{1}{\Phi(h)} \left[\Phi(h) - \frac{h}{\sqrt{2\pi}} e^{-\frac{h^2}{2}} - \frac{1}{2\pi \Phi(h)} e^{-h^2} \right]$$

$$\text{Var}(y|x < h) = m_{02} - m_{01}^2 = \frac{1}{\Phi(h)} \left[\Phi(h) - \frac{\rho^2 h}{\sqrt{2\pi}} e^{-\frac{h^2}{2}} - \frac{\rho^2}{2\pi \Phi(h)} e^{-h^2} \right]$$

$$\begin{aligned} \text{Cov}(x, y|x < h) &= m_{11} - m_{10} m_{01} = \frac{1}{\Phi(h)} \left[\rho \Phi(h) - \frac{\rho h}{\sqrt{2\pi}} e^{-\frac{h^2}{2}} - \frac{\rho}{2\pi \Phi(h)} e^{-h^2} \right] \\ &= \rho \text{Var}(x|x < h) \end{aligned}$$

The conditional correlation is then defined as:


$$\begin{aligned}
 \text{Corr}(x, y|x < h) &= \frac{\text{Cov}(x, y|x < h)}{\sqrt{\text{Var}(x|x < h) \text{Var}(y|x < h)}} \\
 &= \rho \frac{\text{Var}(x|x < h)}{\sqrt{\text{Var}(x|x < h) \text{Var}(y|x < h)}} \\
 &= \rho \sqrt{\frac{\text{Var}(x|x < h)}{\text{Var}(y|x < h)}} = \rho \sqrt{\frac{\Phi(h) - \frac{h}{\sqrt{2\pi}} e^{-\frac{h^2}{2}} - \frac{1}{2\pi} e^{-h^2}}{\Phi(h) - \frac{\rho^2 h}{\sqrt{2\pi}} e^{-\frac{h^2}{2}} - \frac{\rho^2}{2\pi} e^{-h^2}}}
 \end{aligned}$$

APPENDIX B

BITCOIN EXCHANGE TRADE VOLUME (USD)

EXHIBIT B1

Bitcoin Exchange Trade Volume (USD, logarithmic scale) on Major Exchanges (March 2013 through May 2021)




SOURCE: [Blockchain.com](https://blockchain.com).

APPENDIX C

DISTRIBUTION OF MONTHLY BITCOIN RETURNS

EXHIBIT C1

Distribution of Monthly Bitcoin Returns (July 2011 through December 2020)



ACKNOWLEDGMENTS

This material is for informational purposes only. The views expressed in this material are the views of the authors, are provided “as-is” at the time of first publication, are not intended for distribution to any person or entity in any jurisdiction where such distribution or use would be contrary to applicable law, and are not an offer or solicitation to buy or sell securities or any product. The views expressed do not necessarily represent the views of State Street Global Markets®, State Street Corporation® and its affiliates, and/or Windham Capital Management.

REFERENCES

- Ang, A., and J. Chen. 2002. “Asymmetric Correlations of Equity Portfolios.” *Journal of Financial Economics* 63 (3): 443–494.
- Anson, M. 2002. “Managed Pricing and the Rule of Conservatism in Private Equity Portfolios.” *The Journal of Private Equity* 5 (2): 18–30.
- . 2007. “Performance Measurement in Private Equity: Another Look.” *The Journal of Private Equity* 10 (3): 7–22.
- . 2013. “Performance Measurement in Private Equity: Another Look at the Lagged Beta Effect.” *The Journal of Private Equity* 17 (1): 29–44.
- Asness, C., R. Krail, and J. Liew. 2001. “Do Hedge Funds Hedge?” *The Journal of Portfolio Management* 28 (1): 6–19.
- Czaronis, M., M. Kritzman, and D. Turkington. 2021. “The Stock-Bond Correlation.” *The Journal of Portfolio Management* 47 (3): 67–76.
- Kinlaw, W., M. Kritzman, and D. Turkington. 2014. “The Divergence of High- and Low-Frequency Estimation: Causes and Consequences.” *The Journal of Portfolio Management* 40 (5): 156–168.