

NOVEMBER 2022

Applying Machine Learning Portfolio Modeling to Bitcoin

New Modeling Techniques to Evaluate Bitcoin in a Multi-Asset Class Portfolio

Chris Kuiper, CFA, Director of Research, Fidelity Digital Assets

Aaron Gao, Head of Investment Application Development,
Fidelity Advanced Technologies for Investment Management (ATIM)

Yonatan Tekleab, PhD, Director of Data Analytics, Fidelity ATIM





Introduction

Investors considering whether to add bitcoin to their multi-asset class portfolio face the challenge of how to think about both the impact of a bitcoin allocation and how to set and manage their exposure. A simple way to start can be to examine a portfolio by comparing what its historical return would be with and without bitcoin, which is what we have done in “Getting off Zero,”¹ our previous report. However, the major shortcoming of this approach is that bitcoin’s price history is incredibly short, and its historical returns have been abnormally high as the asset class matures (bitcoin’s market capitalization went from a value of less than \$10 million to more than \$1 trillion in approximately 10 years).

In this paper, we explore a different way to understand the dynamic impact that bitcoin may have on a portfolio, using new, proprietary portfolio modeling technology from Fidelity. This technology can be used to support investment decisions and help explore scenarios, but is not meant to give a specific forecast or portfolio recommendations.

- Our new model uses machine learning and high-performance cloud computing to overcome many of the limitations inherent in short time series data, as well as some shortcomings of traditional statistical investment models.
- To enable personalized decision support, the model allows for flexible returns assumptions. Here we look at arguably more realistic return assumptions that bitcoin may achieve going forward, compared with its outsized performance of the past. The characteristics of bitcoin’s return outcomes considering such expected return assumptions are explored in a portfolio context.
- We also extend our previous analysis from a simple 60% equity and 40% bond portfolio to one that is more representative of today’s institutional and multi-asset class portfolio that could include real estate, commodities, precious metals, and cash.
- Finally, we introduce the mathematical concept of evaluating a portfolio’s risk-return “efficient” frontier, defining investment risk not only on a basis of volatility but also extreme loss potential. We advance this “return loss efficient frontier” as more informative for an emerging asset class like bitcoin, in which investors want to take advantage of the potential upside gain but still control overall portfolio loss.

Summary of Findings

Fidelity is increasingly focused on providing new capabilities that analyze digital assets. Our analysis framework provides one way for investors to consider potential return outcomes under different market environments and make informed investment decisions according to their personal risk tolerance.

1. See <https://www.fidelitydigitalassets.com/research-and-insights/getting-zero-bitcoins-role-modern-investment-portfolios>.



Key Observations:

- Bitcoin’s observed return patterns share similarities with other emerging assets, including periods of sharp acceleration and contraction as adoption has expanded.
- Portfolio allocation guidance builds from crucial asset return and investor risk tolerance assumptions. Our framework supports personalizing these assumptions directly or exploring alternatives comparatively.
- Adding bitcoin may lift a portfolio’s risk-reward profile, depending on bitcoin’s expected return assumption relative to other portfolio assets, its correlation with other portfolio assets, and any specified allocation constraints.
- Bitcoin’s volatility and its correlation with other portfolio holdings affect the portfolio’s return outcomes. Examination of full asymmetric return distributions can extend “efficient” frontier methodology to capture bitcoin’s portfolio risk impact more accurately than a traditional statistical risk model, and to inform trade-offs with other holdings.
- Reviewing return outcomes under different investment horizons and market scenarios provides greater insight, given the dynamic nature of bitcoin performance relative to traditional assets.
- Ongoing model calibration will be important to inform investor rebalancing considerations as digital asset markets continue to evolve.

Our modeling infrastructure can measure complex and changing asset relationships to calculate portfolio-level probabilistic outcomes that are also sensitive to market scenarios. The analysis that follows overcomes simplifying assumptions used in traditional statistical models, which may neglect the skewed nature of investment returns, a characteristic of bitcoin that has increased investor interest.

The Starting Point—Historical Analysis

Our starting baseline portfolio is representative of a typical institutional portfolio. We have kept this relatively simple because the point of our analysis is to focus on how adding bitcoin to a representative portfolio will affect the allocation characteristics.

The Baseline Portfolio Allocation

Our baseline portfolio is as follows, using expected returns based on historical data going back to 1990:

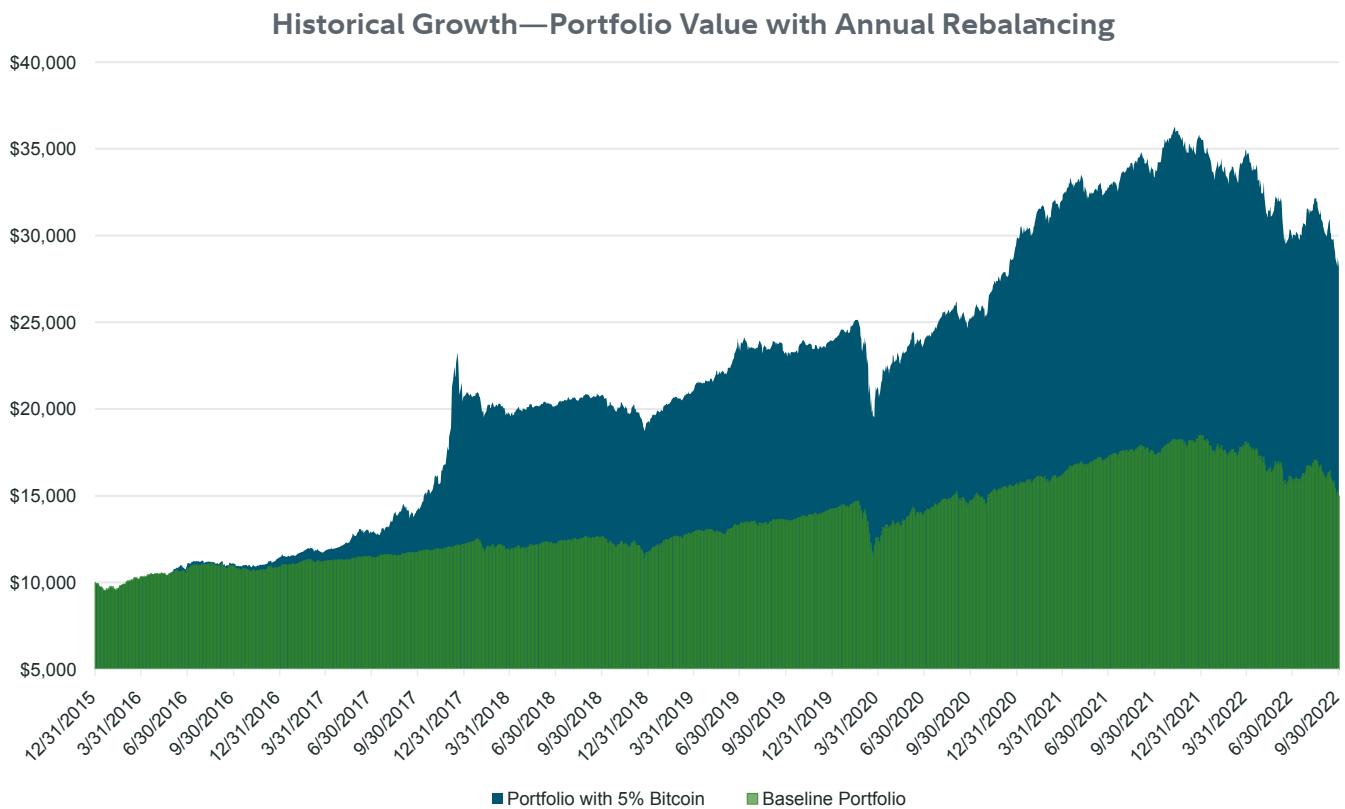
Weight	Asset Class	Instrument	Expected Return
45%	Equities	Fidelity® 500 Index Fund (FXAIX)	12.21%
35%	Fixed Income	Fidelity® U.S. Bond Index Fund (FXNAX)	5.76%
10%	Real Estate	Fidelity® Real Estate Index Fund (FSRNX)	13.75%
5%	Commodities	Fidelity® Series Commodity Strategy Fund (FCSSX)	1.08%
3%	Precious Metals	Fidelity® Select Gold Portfolio (FSAGX)	8.85%
2%	Cash	Fidelity® Government Money Market Fund (SPAXX)	1.84%



Historical Performance with and without Bitcoin

Figure 1 shows the baseline portfolio since January 2016 (selected according to bitcoin’s adoption), with and without bitcoin, assuming annual rebalancing and a starting value of \$10,000. The portfolio with bitcoin assumes a 5% allocation taken from the equity portion of the baseline portfolio (equities reduced to 40% from 45%, and 5% bitcoin added). The period from December 31, 2015, through September 30, 2022, was selected to capture bitcoin’s periods of extreme gain and loss.

Figure 1.



Portfolio (\$10,000 Starting Balance)	Ending Balance	CAGR	Return Volatility (std. dev.)	Sharpe Ratio	Largest Drawdown (max peak-to-trough loss)
Baseline—no bitcoin	\$15,068	6.26%	8.51%	0.74	-\$3,515 (-19%)
Baseline + 5% bitcoin	\$28,384	16.72%	23.05%	0.73	-\$8,053 (-22%)

Historically, if bitcoin were a part of this portfolio, the ending balance would have been \$13,316 higher (+88%), as the compound annual growth rate (CAGR) increased to 16.7% from 6.3%. While volatility also would have increased, this would have been compensated for by the extra return. Despite bitcoin’s large return volatility, a 5% allocation taken from equity over this period results in a nearly similar Sharpe Ratio as that of a portfolio without bitcoin. For assets like bitcoin, upside returns have far exceeded an investment’s 100% loss potential, and volatility may not be the best risk measure. An alternative risk measure is the impact on a portfolio’s worst outcomes. Over this roughly 7 year period, the largest peak-to-trough loss



for the portfolio with bitcoin is greater, but not by enough to offset the overall gain in comparison to the baseline portfolio. This is due to the small allocation to bitcoin and the historically low correlation of bitcoin to the other asset classes. During this period, the baseline portfolio with a 5% bitcoin allocation outperformed the baseline portfolio without a bitcoin allocation.

Introducing Our New Model and Efficient Frontiers

The previous illustration can be a useful starting point, but one of the biggest drawbacks is that it uses only historical data in a very static dimension. This is particularly an issue because bitcoin's performance history is short, extremely volatile, highly skewed toward positive return outcomes, and marked by changing correlations to traditional market factors. Given bitcoin's relatively young age as an investable asset, historical returns are not likely to repeat in the future. Bitcoin's return correlations with other assets are also not likely to stay the same.

We therefore applied a new and advanced portfolio modeling capability from Fidelity that can overcome such shortcomings of historical playback analysis. This capability uses machine learning, including a type of generative artificial intelligence, along with high-performance cloud computing to produce simulated market scenarios and the respective simulated return outcomes for hundreds of thousands of instruments including digital assets. With bitcoin, it is particularly important to capture dynamic asset class correlations and extend information from history to explore developing interdependencies. While respecting asymmetries in historical and simulated return outcomes, this approach enables the examination and evaluation of changes to a portfolio's probable return outcomes under different allocations and across market scenarios. For example, the probable return outcomes in recession markets generally look quite different from those of mid-phase business cycle markets because of changes in correlations among assets under these different economic conditions. Specifically for our purpose in considering bitcoin allocations, we use our model's ability to specify our own return assumptions for bitcoin, and thereby examine much more modest prospective levels than historical past performance.

Overall, the purpose of our model is to inform investment decisions by simulating market environments to provide a range of outcomes that illustrate different risk and return tradeoffs. Observations of historical return patterns are the foundation for the framework. These observations are updated routinely to reflect developing characteristics of bitcoin and its relationships with traditional assets.



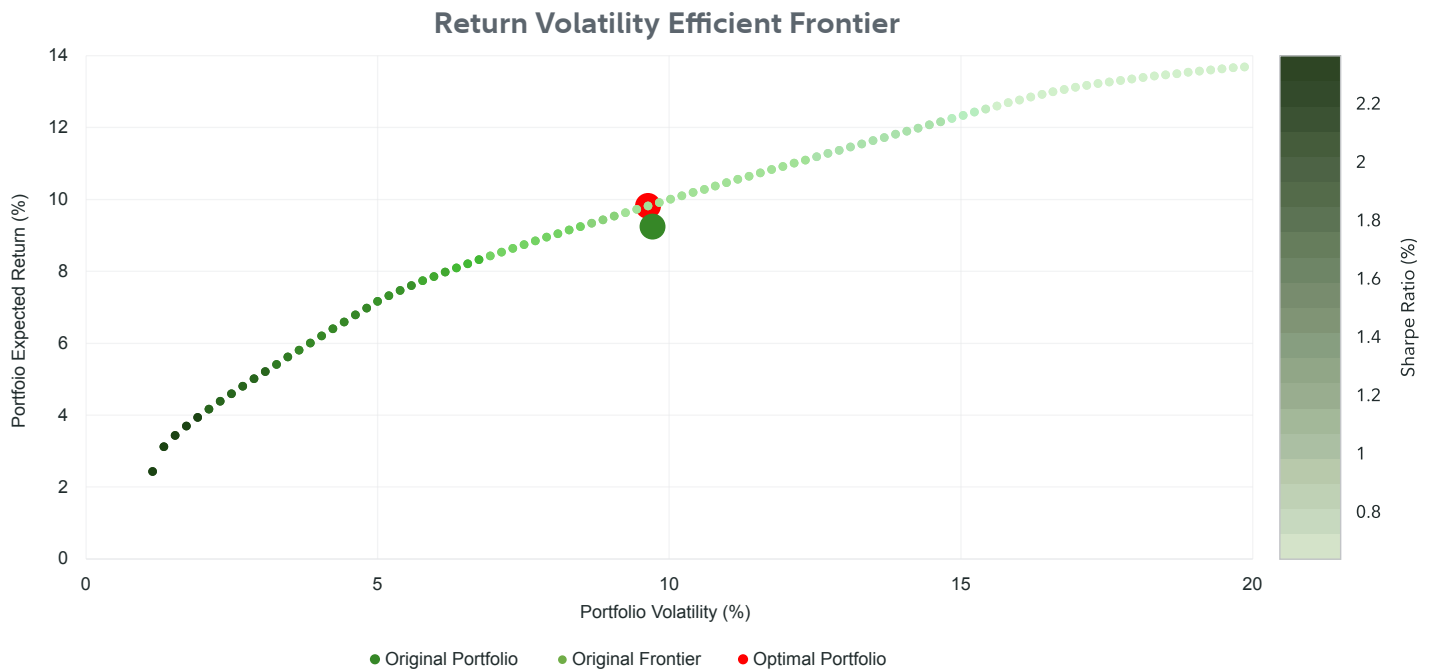
The Baseline Portfolio's Efficient Frontier without Bitcoin

Many professional investors are familiar with the concept of an efficient frontier, a staple of modern portfolio theory that plots the optimal portfolio expected return (the mean of an underlying modeled probability distribution) for a given level of portfolio risk. The y-axis of an efficient frontier graph represents the portfolio expected returns, while the x-axis represents a measure of risk. A frontier plot illustrates the mathematically calculated optimal risk-versus-reward profile given an investable universe, allocation constraints, return assumptions, and market views.

The importance of model methodology and input assumptions should not be overlooked; an efficient frontier does not exist in isolation from its inputs. Similarly, the same inputs under differently described market scenarios will have a different frontier. For example, the efficient frontier for the same investable securities, return assumptions, and allocation constraints, but under a stagflation market scenario, will look quite different from that calculated across all market scenarios.

Our modeling framework extends the use of frontier graphs. Behind a frontier plot, thousands of simulated return outcomes are calculated for selected assets under a selected market scenario while factoring in user preferences. The mean of the simulated returns is displayed as the expected return on the y-axis. For analysis using volatility as the definition of risk, the standard deviation of the simulated returns is represented along the x-axis. Using our model, we can calculate and plot the efficient frontier of our baseline portfolio against the portfolio's original allocation. Figure 2 shows that, given the input assumptions, there is a different weighted mix of assets that has a higher mean return with the same volatility. This specific mix is illustrated by the *optimal portfolio* (●) on the frontier, which lies above the *original* (baseline) *portfolio* (●).

Figure 2.

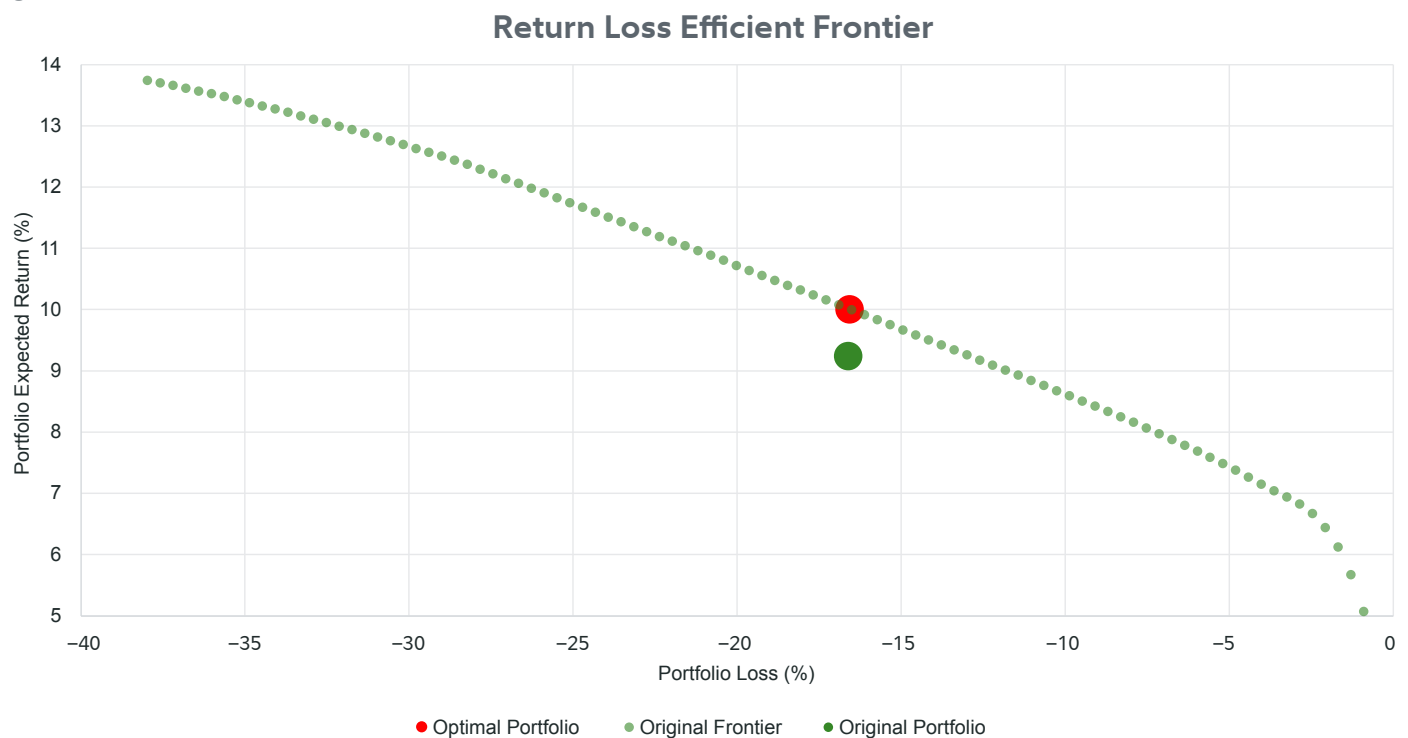




This efficient frontier uses historical return assumptions for each of the assets, a one-year time horizon, and no allocation limits or caps to any of the portfolio constituents.

While many investors are certainly accustomed to viewing and working with this traditional definition of volatility being the risk measure, we present a new frontier, replacing portfolio volatility with the average of the worst 5% portfolio return outcomes (i.e., extreme loss potential), as the measure of risk along the x-axis. In other words, rather than plotting the maximum expected return per unit of volatility, we plot the portfolio's maximum expected return per unit of extreme loss, or the level of pain that can be endured over the next year at a 5% probability. We label this curve the "return loss efficient frontier." This is only possible because our simulation model supports asymmetry of return outcomes as well as dynamic asset dependency structures.

Figure 3.



The reason for preferring this loss-based view is because volatility is the standard deviation of the portfolio's returns, describing both upside and downside volatility. Of course, investors are usually more concerned with downside risk; therefore, a look at loss outcomes can be more insightful, especially once bitcoin is added to the portfolio. Figure 3 shows the return loss efficient frontier for our original (baseline) portfolio. Keeping the portfolio (loss) risk unchanged, we can rebalance to increase our expected return nearly a full percentage point on a one-year horizon.



How Is Bitcoin Modeled?

One of the first challenges investors face when considering and seeking to understand bitcoin is how to classify it. Is it similar to a technology company? A network or platform? A commodity, a currency, or “digital gold?” Is bitcoin even an asset? We have written on this topic in the past,² but for the scope and purposes of this paper, we are going to explore bitcoin from a purely quantitative approach. We begin by examining bitcoin’s return history and characteristics compared with some other modeled assets, as observed in Figures 4 and 5.

Figure 4.


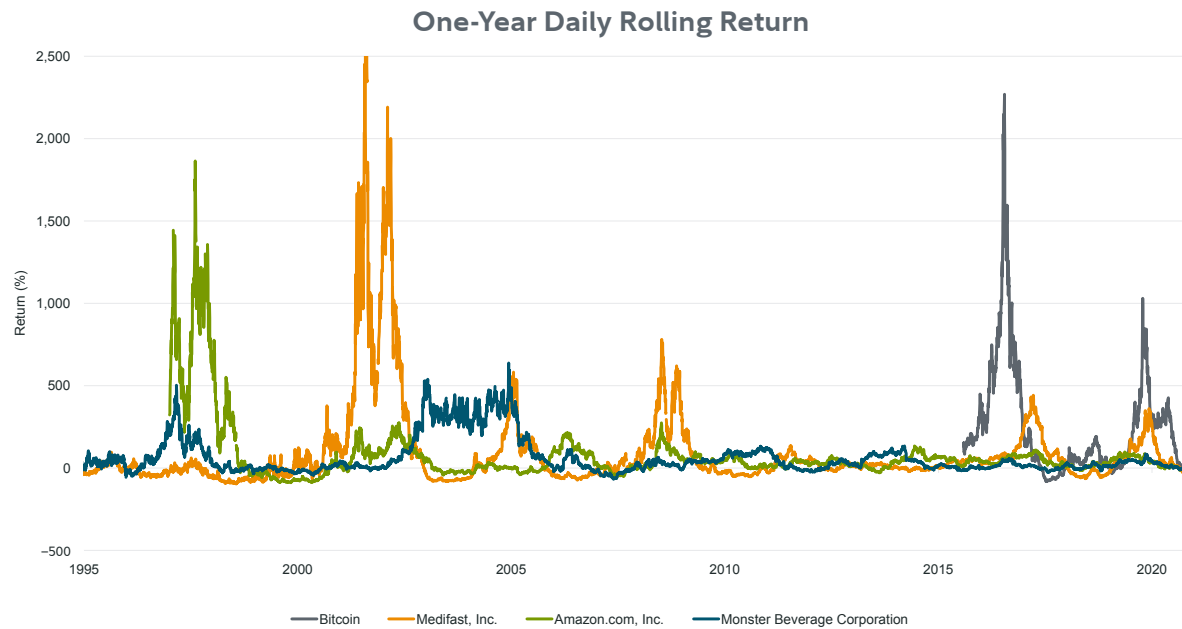


Figure 5.

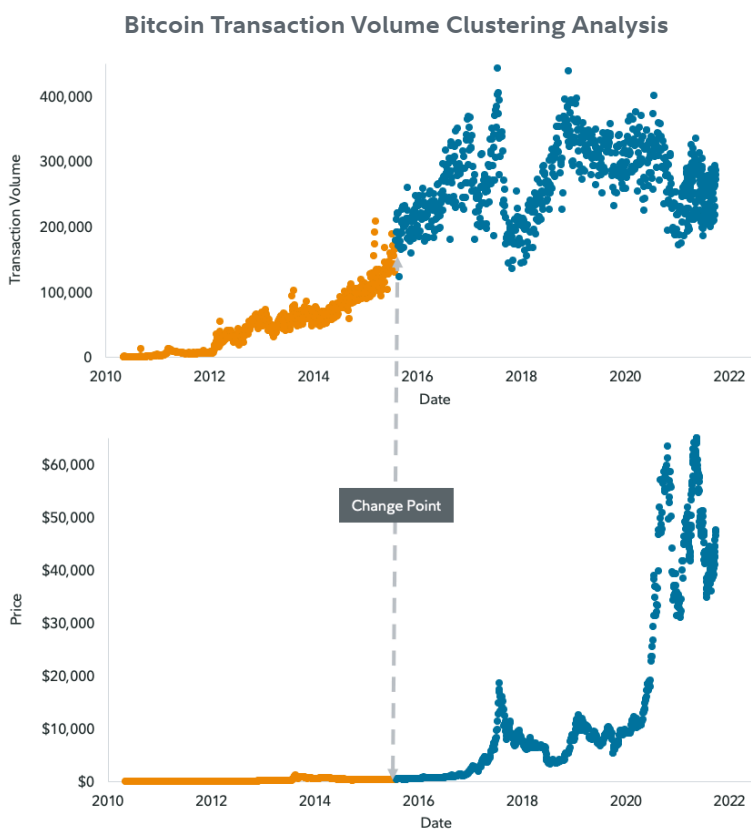


2. See, for example: <https://www.fidelitydigitalassets.com/research-and-insights/bitcoin-investment-thesis-bitcoin-aspirational-store-value-system> and <https://www.fidelitydigitalassets.com/research-and-insights/bitcoin-first>.



In the illustration above, we observe that bitcoin has similarities with other extremely high-growth and high-return equities, such as Medifast, Amazon, and Monster Beverage. Notably, they have all experienced exponential growth and high volatility, with episodes of severe loss. In statistical terms, these assets have highly skewed return distributions, with very long positive tails (occurrences of gains). Although the negative tail (occurrences of losses) is fat, the positive tail is much longer. In many ways, from a purely quantitative perspective, bitcoin has resembled an exaggerated emerging equity instrument, or perhaps a “growth stock on steroids.” These assets also share a common maximum loss potential of 100%. In a well-diversified portfolio, volatility of individual assets is not always the enemy of good risk management. Volatility can be driven by exceptional growth.

Figure 6.



Another factor that needs to be considered in modeling bitcoin is how much of its price history to include. Bitcoin’s early trading was arguably a period of intense price discovery, with thin trading and low volumes on exchanges that were less efficient and transparent as compared with those today.

Therefore, our model uses bitcoin’s realized price history starting from January 2016. A machine learning–identified significant change point, considering bitcoin’s price and transaction volume, is illustrated in Figure 6. Bitcoin’s price can then be imputed through earlier market environments while capturing the realized volatility, correlation, and asymmetric return patterns observed after the change point.

Using generative artificial intelligence, bitcoin’s simulated returns reflect the asset’s volatile history, dynamic correlations to market factors, and observed skewed distribution. Importantly, as our model is routinely recalibrated, future evolution of bitcoin trading patterns will be captured. In essence, the process builds from observed history and anticipates the need to identify characteristic changes.

The Efficient Frontier with Bitcoin Added

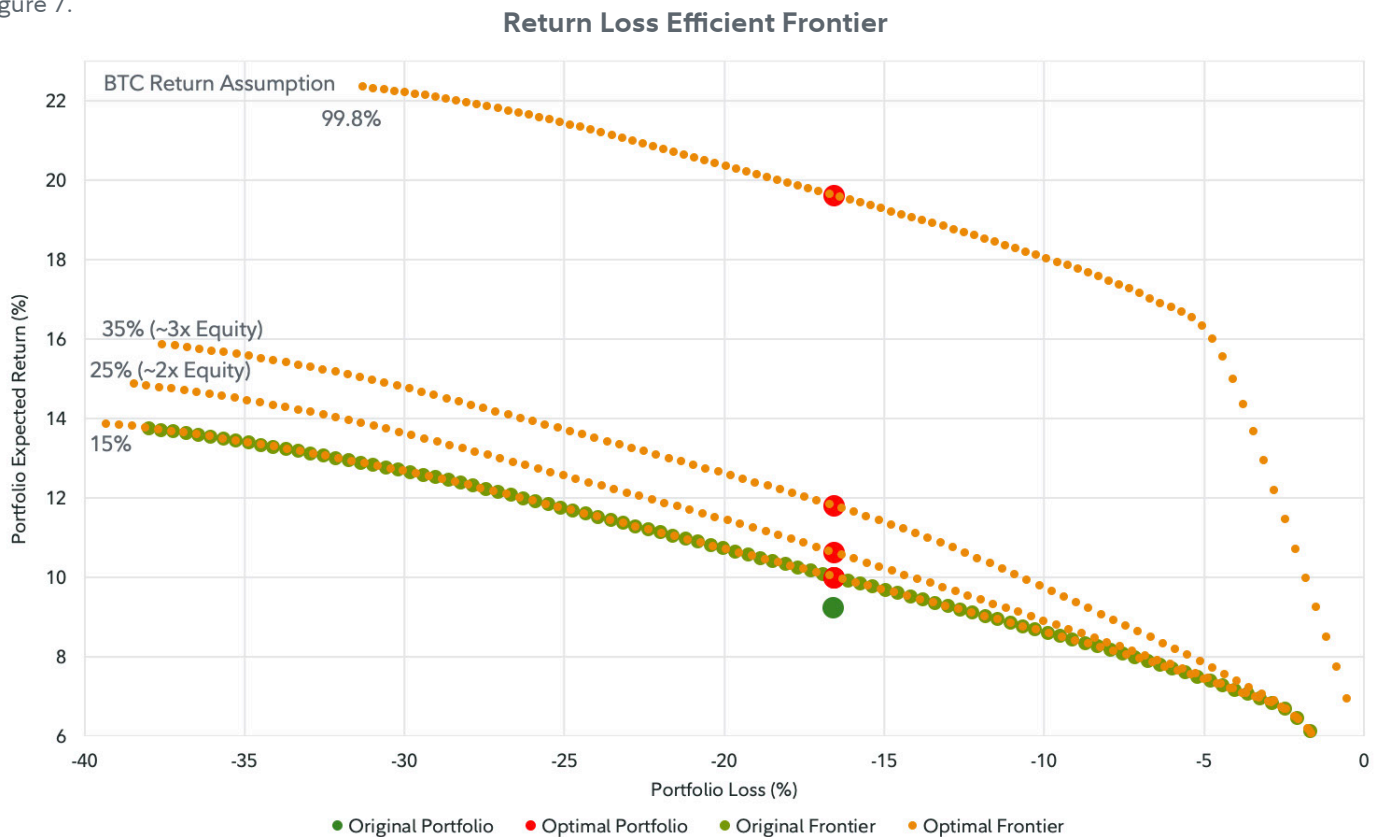
Return assumptions are key in efficient frontier analyses. Historical return assumptions often do not provide good guidance for future returns, unless history is expected to repeat. Still, historical data can provide a reference of past performance that can contextualize our forward outlook. Using bitcoin’s past performance to inform its forward expected return assumption would yield a dramatic rise in the portfolio’s return loss



efficient frontier. This is because the return assumption for the bitcoin model using historical data is 99.8% per year—something bitcoin achieved in the past but is arguably less likely to do in the future. Many investors object to using bitcoin’s historical returns, and rightly so, as an asset worth approximately \$500 billion today growing at a 99.8% CAGR would reach the value of the entire global fixed income market (estimated at \$127 trillion) in only eight years.³

However, with our model, we can explore our own return assumptions for bitcoin, while still maintaining the other characteristics, such as the high volatility, large skew, and potential for large drawdowns. Several example frontiers are shown in Figure 7, with varying return assumptions for bitcoin. The examples that follow use bitcoin possibilities set to align with historical outcomes of traditional portfolio assets. These are not Fidelity forecasts. With lower expectations for all assets, the same relative analysis could be applied.

Figure 7.



If bitcoin is added to our baseline portfolio, using a 15% annualized return assumption and capping the maximum bitcoin allocation at 10%, the return loss efficient frontier overlaps the frontier of the baseline portfolio, providing no lift in the risk-return profile. This is because the other portfolio asset return assumptions are high and a 15% return assumption for bitcoin does not provide enough value to justify bitcoin’s added risk. In other words, if someone does not believe bitcoin will return more than 15% per year, then there would be no reason to add it to the portfolio as it would only introduce more risk with no additional return. A return assumption of 25% improves the profile substantially at more negative loss tolerance levels. Increasing to 35% would yield a greater

3. See <https://www.sifma.org/resources/research/fact-book/>.




profile lift, even at lower loss tolerance levels, by starting bitcoin allocation at very low loss tolerances.

Let's consider the 35% average annual compound return assumption. To demonstrate directionality, we are presenting a seemingly optimistic expected return for bitcoin. To put this in perspective, this would be equivalent to bitcoin's market cap growing to approximately 50% of the market cap for gold over the next eight years (assuming a \$500 billion current market cap for bitcoin and that the market cap for gold remains the same), and if this CAGR continued, bitcoin would reach a market cap of approximately \$10 trillion in the next decade, still lower than the current gold market cap of approximately \$11 trillion. It is not as important to predict the specific annual return for bitcoin here, but rather to input a much lower and more reasonable assumption than what bitcoin has done historically. Granted, this is "more reasonable" on a relative basis compared with bitcoin's high historical performance (nearly 100% CAGR), but still aggressive compared to other, more mature, asset classes. A lower return for bitcoin could therefore be considered by a user, especially if the user is also lowering the returns for other asset classes down from the historical (but currently high) assumptions. In other words, as long as you lower both bitcoin return assumptions and historical asset class return assumptions, the relative relationship holds and the addition of bitcoin may provide a lift to the return profile as long as your bitcoin return assumption is greater than the combined return assumption for traditional assets. Again, if bitcoin's return assumption is equal or less, it will not provide any benefit to a portfolio. In this example, bitcoin's 35% expected return assumption is roughly 3x that of equities. These are not return forecasts by Fidelity and are used only as illustrations.

With this 35% return assumption, keeping all other parameters the same (such as capping bitcoin allocation at 10%), we see the return loss efficient frontier with and without bitcoin added in Figure 8.

Figure 8.





There are two key takeaways to note: First, the frontier has shifted upward, meaning that the addition of bitcoin has indeed increased the portfolio's expected return. This is true at any threshold of portfolio loss exceeding -2.1% , above which the optimized portfolio holds some bitcoin. Second, the bitcoin allocation increases the return potential well above that of the baseline portfolio. Moving from right to left along the frontier, we observe that the separation between the frontiers continues to grow until it reaches the 10% allocation cap for bitcoin at a portfolio extreme loss potential of -12.1% . Beyond this point, the frontiers maintain nearly the same separation.

This model is showing that for a given portfolio loss tolerance, a higher expected return can be achieved by adding bitcoin to the portfolio. Ultimately, a higher expected return that was not possible with the baseline portfolio alone can be achieved.

It should be noted that this model and this paper are NOT making a specific bitcoin allocation recommendation. Our model is being used to explore specific assumptions within the representative baseline portfolio. By showing the potential improvement in expected return for a given level of risk (portfolio volatility or extreme loss) if bitcoin is added, the approach may help investors consider allocation sensitivities.

Beyond Expected Returns—What Happens to Probability Distributions?

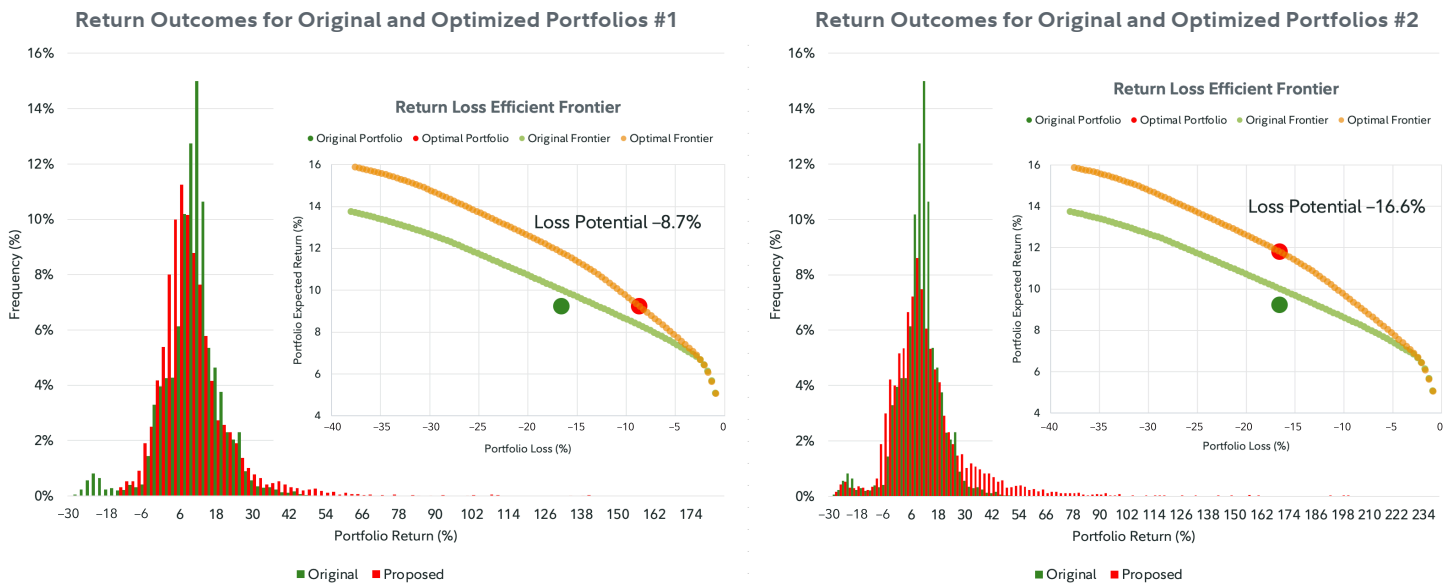
Using our model, we can look beyond the expected returns of optimally weighted portfolios across a range of risk levels. Diving deeper into any of the optimized portfolios (represented by a single point on the frontier) allows us to examine the full range of its potential return outcomes. We can see that adding bitcoin affects the entire distribution of the portfolio in interesting ways.

When an extremely volatile asset's downside is bounded but its upside is unbounded, it does not require a large allocation to enable participation in that upside. In a portfolio context, what becomes important are the broader tradeoffs. Given flexibility to add up to 10% bitcoin, the optimal portfolio will take from equity to fully fund that allocation and will move additional equity as well as other asset allocations to fixed income to take advantage of its lower loss potential and weak correlation to bitcoin.

In doing so, the optimized portfolio with bitcoin increased upside potential by extending the positive tail. Using bitcoin to increase upside potential, the optimizer removed equity allocation, thereby reducing the probability of extreme loss. The moderate allocation to bitcoin, however, increased the probability of moderate loss. This becomes more intuitive when we look into the optimal portfolio allocation changes across the frontier.



Figure 9.



		Portfolio	Original	Optimized #1	Optimized #2
Asset Allocation (%)		Equities	45.00	5.33	15.42
		Fixed Income	35.00	73.74	48.01
		Real Estate	10.00	14.00	26.57
		Commodities	5.00	0.00	0.00
		Precious Metals	3.00	0.00	0.00
		Cash	2.00	0.00	0.00
		Bitcoin	0.00	6.92	10.00
Portfolio metrics (%)	Return		9.24	9.25	11.80
	Loss potential		16.61	8.69	16.61

Figure 9 shows two sets of return distributions overlaid on the original baseline portfolio distribution, one in which the portfolio is optimized while maintaining the baseline (original) portfolio's expected return (#1), and the second in which the portfolio is optimized while maintaining the baseline portfolio's loss potential (#2).

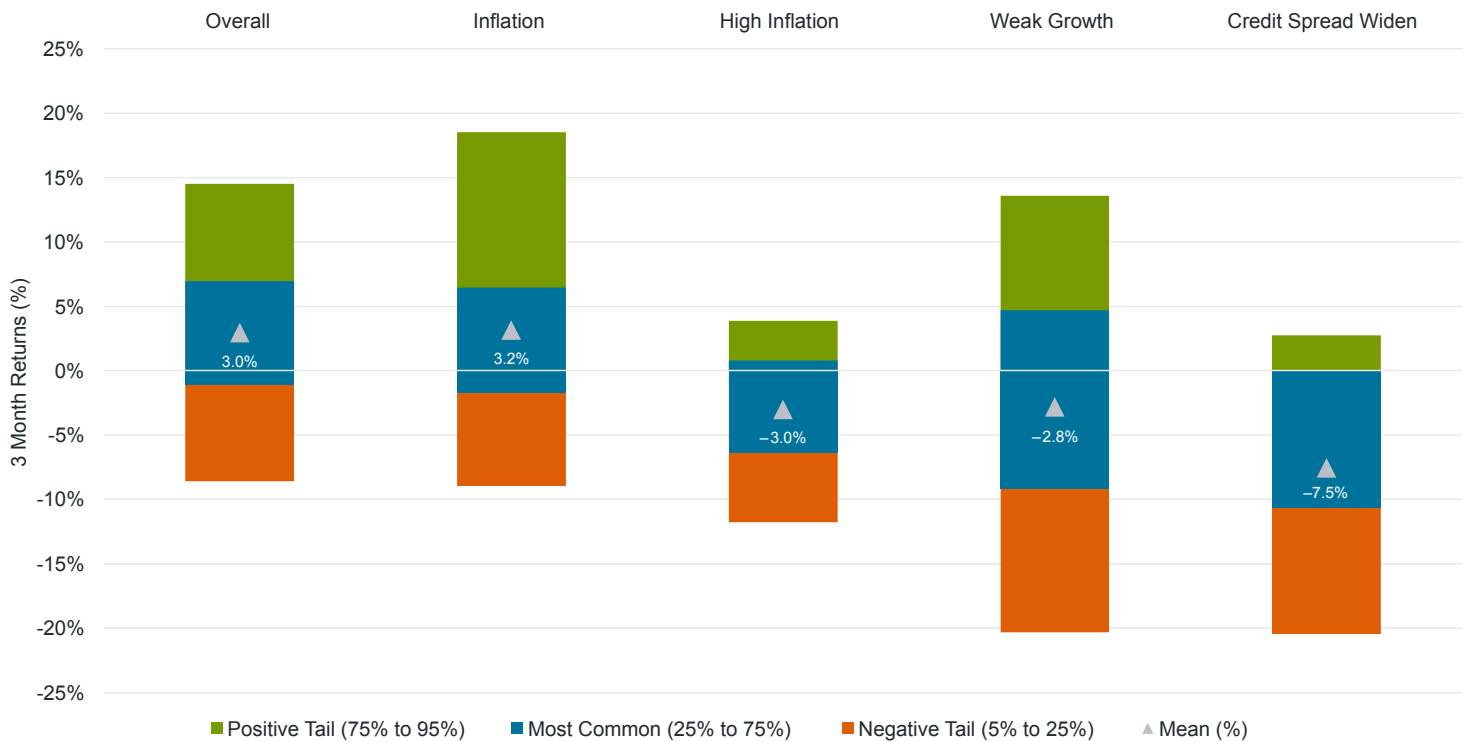
Unbounded reallocation to demonstrate directional change.

Again, our model's ability to capture asymmetric return outcomes helps demonstrate the potential tradeoffs made in a portfolio optimization calculation. These tradeoffs are illustrated in Figure 9, which at first glance appears to be an odd portfolio that is not likely to be used by a typical institutional investor (such as only a 5% allocation to equities in the first optimization). This is intentional as it is an unrestricted optimization meant to illustrate the inherent tradeoffs being made when bitcoin is added. As such, it is not meant to be a representation of a well-balanced portfolio, but an exercise that illustrates the somewhat extreme characteristics of bitcoin of which investors need to be aware. Note that Figure 9 above uses the same assumptions as the previous example (Figure 8), with a 35% annual return for bitcoin, all other asset classes using historical returns, and the bitcoin allocation capped at 10%.

Our model also provides the capability of examining return outcomes under various defined market scenarios. For example, investors may find it helpful to view the possible outcomes for a portfolio with bitcoin in an environment with weak growth or inflation market conditions. These return outcomes are illustrated as comparative box plots in Figure 10 for "Optimized portfolio #2" (shown above where the allocation includes the maximum 10% bitcoin, plus fixed income, equities, and real estate).



Figure 10.



Depending on an investor’s market outlook, their portfolio return outcomes may look quite different. The efficient frontier will change according to the market environment; therefore, the optimal portfolio allocation will also be different under different market scenarios. It is beyond the scope of this paper to describe the details of the market scenarios in Figure 10. Scenario definitions and the length of data included in the analysis will impact scenario outcomes. This figure is for illustration purposes only.

Conclusion

We started with a simple historical analysis in which we did a replay of a standard multi-asset class portfolio with and without bitcoin. Not surprisingly, the portfolio with bitcoin showed much better results, given bitcoin’s meteoric rise, which far outweighed its volatility. While useful in some ways, this rear-view analysis leaves investors wondering what to consider going forward as bitcoin matures into a larger asset class with lower returns more likely.

We then introduced our portfolio modeling capability that takes historical price behavior data but extends it with proprietary machine learning methods. By simulating potential outcomes, we may consider portfolio allocation questions with bitcoin alongside traditional investment assets. Our model analyzes the consistency of trading characteristics to determine relevant historical price data for bitcoin and other emerging assets.



This new modeling capability also allows us to specify our own return assumptions for bitcoin looking forward, which we lowered significantly from its historical average.

Provided the investor is comfortable with the assumptions and the aforementioned asset allocation tradeoff considerations, our model illustrated that the potential expected returns for our baseline portfolio could indeed be improved by adding bitcoin (which was capped at a maximum allocation of 10%), without taking on more volatility or extreme loss potential. Furthermore, our model demonstrated the tradeoff of higher expected returns and fewer events of extreme loss, which can come at the cost of increasing the frequency of more moderate losses.

The information herein was prepared by Fidelity Digital Asset Services, LLC and Fidelity Digital Assets, Ltd. It is for informational purposes only and is not intended to constitute a recommendation, investment advice of any kind, or an offer or the solicitation of an offer to buy or sell securities or other assets. Please perform your own research and consult a qualified advisor to see if digital assets are an appropriate investment option. Price history used for modeling all assets was captured through March 31, 2022.

Custody and trading of digital assets provided by Fidelity Digital Asset Services, LLC, a New York State-chartered, limited liability trust company (NMLS ID 1773897) or Fidelity Digital Assets, Ltd. Fidelity Digital Assets, Ltd., is registered with the U.K. Financial Conduct Authority for certain cryptoasset activities under the Money Laundering, Terrorist Financing and Transfer of Funds (Information on the Payer) Regulations 2017. The Financial Ombudsman Service and the Financial Services Compensation Scheme do not apply to the cryptoasset activities carried on by Fidelity Digital Assets, Ltd.

This information is not intended for distribution to, or use by, any person or entity in any jurisdiction or country where such distribution or use would be contrary to local law or regulation. Persons accessing this information are required to inform themselves about and observe such restrictions.

Digital assets are speculative and highly volatile, can become illiquid at any time, and are for investors with a high risk tolerance. Investors in digital assets could lose the entire value of their investment.

Fidelity Digital Asset Services LLC, and Fidelity Digital Assets Ltd. do not provide tax, legal, investment, or accounting advice. This material is not intended to provide, and should not be relied on, for tax, legal, or accounting advice. Tax laws and regulations are complex and subject to change. You should consult your own tax, legal, and accounting advisors before engaging in any transaction.

Some of this information is forward-looking and is subject to change. Past performance is no guarantee of future results. Investment results cannot be predicted or projected.

Fidelity Digital Assets and the Fidelity Digital Assets logo are service marks of FMR LLC.

© 2022 FMR LLC. All rights reserved.

1044199.1.0