Simple and Effective Portfolio Construction with Crypto Assets

Kasper Johansson Stephen Boyd

November 16, 2024

Abstract

We consider the problem of constructing a portfolio that combines traditional financial assets with crypto assets. We show that despite the documented attributes of crypto assets, such as high volatility, heavy tails, excess kurtosis, and skewness, a simple extension of traditional risk allocation provides robust solutions for integrating these emerging assets into broader investment strategies. Examination of the risk allocation holdings suggests an even simpler method, analogous to the traditional 60/40 stocks/bonds allocation, involving a fixed allocation to crypto and traditional assets, dynamically diluted with cash to achieve a target risk level.

Contents

1 Introduction 3

2 Related work 6
- 2.1 Stylized facts of financial return data . . . 6
- 2.2 Portfolio construction . . . 6
- 2.3 Crypto trading . . . 8

3 Constrained risk allocation 9
- 3.1 Constrained risk allocation problem . . . 9
- 3.2 Solution via convex optimization . . . 10

4 CRA results 11
- 4.1 Data and experimental setup . . . 11
- 4.2 Metrics . . . 13
- 4.3 Results . . . 15
- 4.4 Shapley attributions . . . 19

5 Dynamically diluted 90/10 portfolio 19
- 5.1 DD90/10 results . . . 20

6 Conclusions 20

Figure 1: Normalized prices of BTC, ETH, and SP500.

Table 1: Performance metrics for BTC, ETH, and SP500.

Metric	BTC	ETH	SP500
Return (%)	43.3	47.0	13.7
Volatility (%)	58.1	71.6	19.5
Sharpe	0.75	0.66	0.70
Drawdown (%)	83.3	93.9	33.9

1 Introduction

Since the introduction of cryptocurrencies in 2009 [Nak08], the field of crypto trading has rapidly grown. In this note we consider the problem of constructing a portfolio of assets, including a combination of traditional financial assets such as stocks and bonds with crypto assets such as Bitcoin and Ethereum. Our conclusion is that despite the well known extreme swings of crypto currency values, simple standard portfolio construction methods suffice to realize the benefits of including crypto assets in a portfolio.

Figure 1 shows the normalized prices of Bitcoin (BTC), Ethereum (ETH), and the S&P 500 index (SP500) over the past six years. To the eye it reasonably seems that crypto returns are quite different from traditional returns. Table 1 lists some metrics for these asset returns over the same period, which also suggest that crypto asset returns fundamentally differ from traditional asset returns. For example, ETH at one point dropped in value by a factor of almost 20×, while the maximum drop in value of SP500 is only one third.

Stylized facts of crypto returns, such as extreme volatility, heavy tails, excess kurtosis, and skewness, have been well documented [LT21, HPR19, KRNY24]. Figure 2 shows quantile-quantile (QQ) plots of the log returns of BTC, ETH, and SP500 over the last six years. All three return distributions have tails that deviate significantly from the normal distribution, with the crypto asset returns exhibiting more extreme tail behavior than the market index. The 1st and 99th return percentiles are shown in the plots as dashed green lines; here too we see that crypto asset returns have considerably bigger tails than the market index, even when normalized to have the same volatility.

The documented characteristics of crypto asset returns have made conventional investors hesitant to include them in their portfolios, due to perceptions of risk and unpredictability. While asset managers may have additional concerns, such as legislative risk and other fundamental factors like the ongoing debate about the legitimacy of crypto as a real asset [Yar13, Kru13, Kru11, CK13], we will focus here on the concerns related to the statistical properties of crypto asset returns.

These characteristics of crypto asset returns have also led to a wealth of research on how to extend traditional portfolio construction methods to include crypto assets, including methods similar to those used for traditional assets [Hol20, BM19, Bur19, HRF19], as well as more complex machine learning driven methods [RRdSJB22, JL17, LB20, JXL17, Ram21, RRdSJB22].

Despite the documented differences between crypto and conventional asset returns, some authors have argued that the two asset classes are fundamentally similar [Pal24, Chap. 2], even though the crypto asset returns and volatilities are much higher. We agree with this perspective. In this note we show a simple method for constructing a portfolio of traditional and crypto assets using a risk allocation framework, (hopefully) debunking the idea that novel and complex machine learning approaches are necessary to manage a portfolio that includes crypto assets. Based on a back-test of the risk allocation method, we propose an even simpler portfolio construction method, reminiscent of the traditional 60/40 stocks/bonds split, which consists of a 90/10 split of traditional and crypto assets, followed by dynamic (time-varying) dilution with cash, to achieve a given ex-ante risk. We refer to this simple portfolio as DD90/10.

Outline. In §2 we review previous and related work. We introduce in §3 an approach for constructing a portfolio that combines traditional and crypto assets within a risk allocation framework. We illustrate the performance of this method on historical data in §4. Finally, in §5 we propose the DD90/10 portfolio allocation strategy, and show that it has performance similar to the risk allocation strategy.

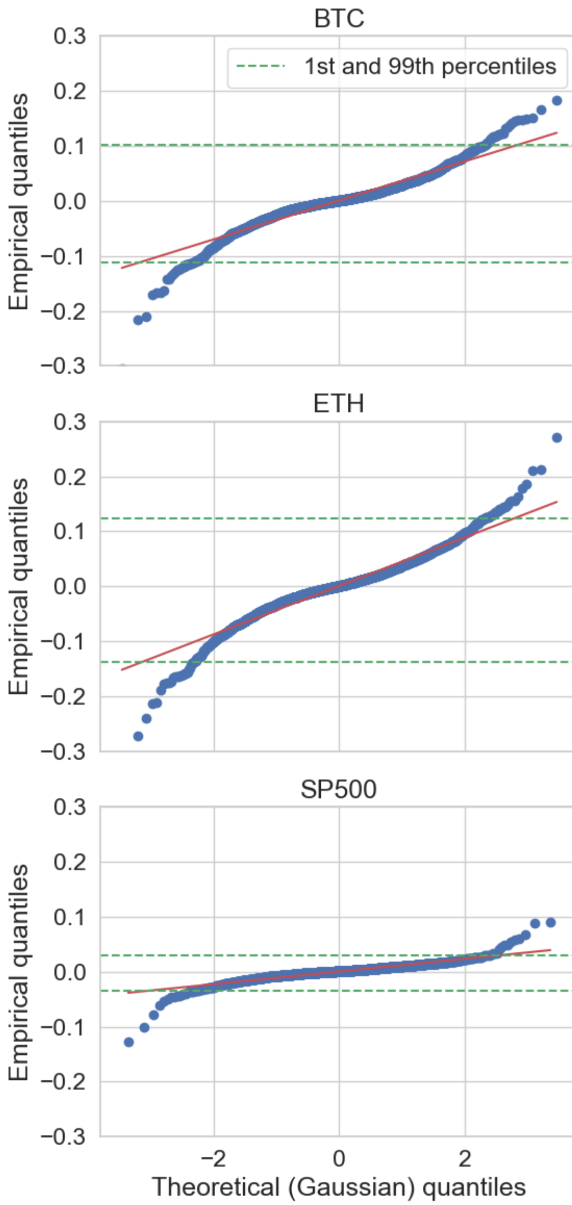

Figure 2: Quantile-quantile plot of log returns of BTC, ETH, and SP500.

2 Related work

2.1 Stylized facts of financial return data

Stylized facts of financial return series refer to a set of empirical observations that are consistently observed across different financial markets and asset classes. We review some of the most prominent stylized facts for financial assets, and refer the reader to [Pal24, Chap. 2] for a more comprehensive overview.

Equities. Stylized facts for equity returns include the non-normal distribution of returns, characterized by heavy tails and excess kurtosis [MT10, BB06]. Another fact is volatility clustering, where large price movements tend to be followed by additional large movements, regardless of direction. While returns themselves are generally uncorrelated over time, the absolute or squared returns often display strong autocorrelation, highlighting a pattern in the magnitude of fluctuations. Furthermore, the leverage effect is a notable feature, where negative returns increase future volatility to a greater extent than positive returns of the same size. Equity return distributions have also been shown to exhibit asymmetry between positive and negative returns [JWZZ20].

Cryptocurrencies. Cryptocurrencies exhibit several patterns similar to traditional financial assets, but with more extreme behavior [HPR19, GBWZ23]. They are highly volatile, with large price fluctuations and heavy-tailed return distributions. Volatility tends to cluster, with periods of high volatility followed by more volatility. While returns show no autocorrelation in the short term, return magnitudes do exhibit autocorrelation. Cryptocurrencies also exhibit asymmetry in returns [KRNY24].

2.2 Portfolio construction

Portfolio construction involves selecting a combination of assets by balancing the trade-off between expected return and portfolio risk. Here we discuss key historical contributions to the field, along with some recent advancements. For more extensive overviews, refer to texts such as [GK00, Chap. 14], [Nar24, Chap. 6], and the studies by [CT06, KTF14, CK23, GB23, BJK+24].

Markowitz portfolio construction. Before Markowitz's seminal work in 1952, portfolio construction was largely based on heuristics and rules of thumb. Markowitz introduced a quantitative framework for portfolio construction, where the return of a portfolio is modeled as a random variable, and the expected return is maximized for a desired level of risk [Mar52]. Despite its simplicity, and having over 70 years of history, the Markowitz model remains the foundation of quantitative investing to this day [BJK+24].

Extensions of the Markowitz model include the Black-Litterman model [BL90, BL92, Bla89], fully flexible views [Meu10], and conditional value-at-risk (CVaR) optimization [RU00], to name a few. The Black-Litterman model is a Bayesian approach to portfolio construction, where the prior distribution is based on the equilibrium market portfolio, and the posterior distribution is updated with user-specific views on the expected returns of the assets. Fully flexible views is a generalization of Black-Litterman, allowing for nonlinear views on the returns. CVaR optimization replaces the variance in the Markowitz model with the conditional value-at-risk of the portfolio, penalizing the tail risk of the portfolio, and directly addresses the issue of a non-normal return distribution.

Machine learning based portfolio construction. Typically, portfolio construction is split into two parts: data modelling and portfolio optimization [Pal24, Chap. 1]. Data modelling is concerned with predicting the expected returns and covariances of the assets, and portfolio optimization concerns selecting a portfolio of assets that trades off expected return and risk. However, with the growing popularity of machine learning, and documented criticism of the Markowitz model [Mic89], in recent years several studies have proposed machine learning based portfolio construction methods as alternatives to this traditional framework [Gur22, BBDRM21]. In particular, it has become popular to combine the two parts of portfolio construction into a single end-to-end machine learning model, where market features are fed into the model, and the model outputs a trade list. The argument for this has been that splitting portfolio construction into two parts is suboptimal, as there are uncertainties in the return forecasts that are not accounted for in the portfolio optimization component of portfolio construction [KX23, §5]. Although theoretically appealing, we have yet to see a wide-spread adoption of these methods in practice, and the Markowitz model remains the dominant framework for portfolio construction [BJK+24]. We refer the reader to [KX23, §5] for a detailed review of end-to-end machine learning models for portfolio construction.

Risk-based portfolio construction. Risk-based portfolio construction methods rely on estimates of the asset return covariances, but do not require an estimate of the expected returns of the assets, making them attractive for practitioners who do not have access to data sources for estimating expected returns reliably. A trivial example of a risk-based portfolio construction method is the equally weighted portfolio. Another popular portfolio is the minimum variance portfolio, which is also the mean-variance efficient portfolio when expected returns are equal. Other risk-based portfolio construction methods include risk parity [Qia11], where the risk contribution of each asset is equal, and maximum diversification portfolios [CC08]. These portfolios can all be computed via convex optimization [JOP+23, §4.4], which makes them reliable, fast, and practical [BV04]. These portfolio construction methods can be implemented in just a few lines of domain specific languages for convex optimization such as CVXPY [DB16].

Risk and covariance estimation. Risk-based portfolio construction methods rely on estimates of the portfolio risk. There are in general two ways to estimate portfolio risk. The first is to use a realized measure of the variance of the portfolio. There are many such methods, including the exponentially weighted moving average (EWMA), methods based on mean absolute deviation or the rolling median [Gea36, Gea47], as well as autoregressive conditional heteroskedasticity (ARCH) and generalized ARCH (GARCH) models [Eng82, Bol86, EB86]. The second way to estimate portfolio risk is to leverage a covariance matrix of the asset returns. The literature on covariance estimating in finance is vast, but probably the most popular method is to use an iterated covariance matrix [Eng02, BB22]. This method decomposes the covariance Σ as $\Sigma = VRV$, where V is a diagonal matrix with the asset standard deviations on the diagonal and R is the correlation matrix. Typically, V and R are estimated separately, using EWMAs with different half-lives. We will refer to this method as the iterated EWMA (IEWMA) [JOP+23, §2.5]. It is also possible to dynamically adjust the half-lives of the EWMAs, to account for time-varying market conditions [JOP+23].

2.3 Crypto trading

Several methods to managing portfolios of cryptocurrencies have been proposed, and these tend to be separated from investment strategies proposed for traditional assets.

Diversification. Many studies have noted that the returns of cryptocurrencies are uncorrelated with traditional assets [BRAMEK21, Yer24, Hol20]. This means that they can be leveraged to diversify a portfolio of traditional assets, and thus increase the risk-adjusted return of the portfolio [Hol20, BRAMEK21].

Portfolio construction. Much of the literature on cryptocurrency trading is focused on machine learning or deep learning [RRdSJB22]. In [JL17] an end-to-end convolutional neural network, taking in raw price data and outputting a trade list, is proposed. A deep Q-learning portfolio management framework is proposed in [LB20]. The authors of [JXL17] introduce a financial-model-free reinforcement learning framework, incorporating convolutional neural networks, recurrent neural networks, and long short-term memory models. In [Ram21] ARIMA models, convolutional neural networks, and long short-term memory methods are used for cryptocurrency price forecasting and multiple portfolio construction strategies are evaluated with the forecasted prices as input. For a more detailed review of machine learning in cryptocurrency portfolio management, see *e.g.*, [RRdSJB22, §4.2].

Some crypto studies use traditional portfolio construction methods, such as the Markowitz framework. In [Hol20], the authors show that crypto assets can improve the performance of a mean-variance optimized portfolio. The authors of [BM19] show how to implement a mean-variance optimized portfolio of cryptocurrencies and that it outperforms an equally weighted portfolio as well as single cryptocurrencies. The paper [GPP20] proposes an extension of the Markowitz model, combining random matrix theory and network measures to manage a portfolio of crypto assets. In [PU19] the authors design crypto portfolios using variance-based constraints in the Black-Litterman model, to account for estimation uncertainties.

Other studies focus on risk-based portfolio construction methods, *i.e.*, those that do not rely on an estimate of the expected returns of the assets. In [Bur19] multiple risk-based portfolio construction methods are evaluated, and the authors find that most of them outperform single cryptocurrencies and the equally weighted portfolio. The authors of [HRF19] find that minimizing the variance and conditional value-at-risk of a portfolio of cryptocurrencies, yields a portfolio that outperforms the market.

3 Constrained risk allocation

The most common approach to portfolio construction is to formulate the problem as a trade-off between expected return and risk, as suggested by Markowitz in 1952 [Mar52]. The main practical challenge with this framework is that it requires estimating expected returns of the assets. These are very difficult to estimate and, for obvious reasons, successful estimation techniques are proprietary; large hedge funds and asset managers have entire teams dedicated to estimating expected returns, using data sources for which they pay large premiums [Pal24]. Here we describe a risk allocation approach, which does not require an estimate of the expected returns of the assets. It is based on the idea of risk parity [Qia11], with additional constraints on the portfolio weights and a risk limit. We refer to this method as constrained risk allocation (CRA).

3.1 Constrained risk allocation problem

We consider a portfolio of n non-cash assets, plus cash. We denote the asset weights as $w \in \mathbf{R}^n$, with $w \geq 0$, where w_i is the fraction of the total portfolio value held in asset i. We let $c \geq 0$ denote the fraction of the portfolio value held in cash, so we have $\mathbf{1}^T w + c = 1$, where $\mathbf{1}$ is the vector with all entries one. We refer to $\mathbf{1}^T w$ as our asset exposure.

Let Σ denote the $n \times n$ (estimate of the) covariance matrix of the returns. Assuming cash is risk-free, the portfolio risk is $w^T \Sigma w$. (The volatility is the squareroot of this.) Inspired by the identity

$$w^T \Sigma w = \sum_{i=1}^{n} w_i (\Sigma w)_i,$$

we define the risk contribution of asset i as $w_i(\Sigma w)_i$. In risk allocation we specify the fraction of risk to be held in each asset, as the vector $\rho \in \mathbf{R}^n$, with $\rho > 0$ and $\mathbf{1}^T \rho = 1$. We interpret ρ_i as the fraction of total portfolio risk contributed by asset i. (We assume that all entries of ρ are positive; if any were zero, we simply would not include that asset in the portfolio.) Thus we have

$$w_i (\Sigma w)_i = \rho_i w^T \Sigma w, \quad i = 1, \ldots, n. \tag{1}$$

The special case $\rho = (1/n)\mathbf{1}$ corresponds to risk parity, where all assets contribute an equal fraction of the risk.

The CRA problem is defined as

$$\begin{array}{ll} \text{minimize} & c \\ \text{subject to} & w \geq 0, \quad c \geq 0, \quad \mathbf{1}^T w + c = 1 \\ & w_i (\Sigma w)_i = \rho_i w^T \Sigma w, \quad i = 1, \ldots, n \\ & w^T \Sigma w \leq \sigma^2, \quad F w \leq g, \end{array} \tag{2}$$

where $w \in \mathbf{R}^n$ and $c \in \mathbf{R}$ are the variables, σ^2 is the maximum allowed risk, and $F \in \mathbf{R}^{m \times n}$ and $g \in \mathbf{R}^m$ describe constraints on the portfolio. In words: We choose the portfolio to minimize cash holdings (or equivalently, maximize asset exposure), subject to a given risk allocation, a total risk limit, and some additional constraints on the weights.

We will assume that $g > 0$ and $F \geq 0$ (elementwise), with each row of F nonzero. We will soon see that this implies the CRA problem (2) always has a unique solution. The weight constraints can be used to enforce a maximum weight on each asset, or a maximum weight on a subset of assets, *e.g.*, crypto assets.

3.2 Solution via convex optimization

The CRA problem (2) is not itself a convex optimization problem, but it can be solved efficiently via convex optimization. We first consider the risk allocation constraints (1) alone, together with $w \geq 0$. It can be shown that w satisfies these constraints if and only if it has the form

$$w = \alpha x^\star,$$

where $\alpha \geq 0$ and $x^\star \in \mathbf{R}^n$ is the unique solution of the convex optimization problem

$$\text{minimize} \quad (1/2)x^T \Sigma x - \sum_{i=1}^n \rho_i \log x_i, \tag{3}$$

with variable $x \in \mathbf{R}^n$ (and implicit constraint $x > 0$). See, *e.g.*, [BV24, §17]. Thus the set of weights that satisfy the risk allocation constraints is a ray, with a direction that can be found by solving a convex optimization problem.

Now we take this very specific form for w and substitute it back into the original CRA problem (2), dropping the risk allocation constraints and $w \geq 0$ since they are automatically satisfied. This gives us the problem

$$\begin{array}{ll} \text{minimize} & c \\ \text{subject to} & \alpha \geq 0, \quad c \geq 0, \quad \alpha \mathbf{1}^T x^\star + c = 1 \\ & \alpha^2 (x^\star)^T \Sigma x^\star \leq \sigma^2, \quad \alpha F x^\star \leq g, \end{array}$$

with scalar variables α and c. Note that the quantities $\mathbf{1}^T x^\star$, $(x^\star)^T \Sigma x^\star$, and $F x^\star$ are constants in this problem.

We can solve this simple problem analytically. Minimizing c is the same as maximizing α. Along with $\alpha \geq 0$, all constraints on α are (positive) upper limits:

$$\alpha \leq \frac{1}{\mathbf{1}^T x^\star}, \quad \alpha \leq \frac{\sigma}{((x^\star)^T \Sigma x^\star)^{1/2}}, \quad \alpha \leq \frac{g_i}{(F x^\star)_i}, \quad i = 1, \ldots, m.$$

(Each of the denominators is positive due to our assumptions and $x^\star > 0$.) It follows that the solution is

$$\alpha^\star = \min \left\{ \frac{1}{\mathbf{1}^T x^\star}, \frac{\sigma}{((x^\star)^T \Sigma x^\star)^{1/2}}, \frac{g_i}{(F x^\star)_i}, \quad i = 1, \ldots, m \right\}. \tag{4}$$

Roughly speaking: Scale the unconstrained risk allocation weights as large as possible with all constraints holding.

Summary. The two step solution procedure is summarized as follows:

1. Solve the optimization problem (3) to obtain $x^\star$.
2. The unique solution of the CRA problem is then given by $w^\star = \alpha^\star x^\star$ where $\alpha^\star$ is given by (4).

We note that Feng and Palomar have suggested a more sophisticated formulation of the CRA problem, in which the risk allocations need only hold approximately [FP15, WFP20], [Pal24, Chap. 11]. This formulation can be approximately solved by solving a sequence of convex problems. Our simple formulation is, however, good enough for us to make our main point.

Variation. We can modify the way $\alpha^\star$ is computed in (4). Instead of estimating the standard deviation of the unconstrained risk allocation portfolio $x^\star$ as $\left((x^\star)^T \Sigma x^\star\right)^{1/2}$, we compute the realized return trajectory of the portfolio $x^\star$, *i.e.*, $(x^\star)^T r_t$, $t = 1, 2, \ldots$, where r_t is the vector of asset returns at time t. We then compute an estimate of the standard deviation of the portfolio return trajectory, using, *e.g.*, a EWMA. Thus to compute the scaling (which sets the cash dilution) we directly estimate the standard deviation of the return trajectory with the unconstrained risk allocation weights, rather than find it from our estimated covariance matrix (which we use to compute the risk allocation weights $x^\star$ in step 1.) We have found that leads to a modest but significant improvement in portfolio performance.

4 CRA results

4.1 Data and experimental setup

Data. We consider daily close prices of two crypto assets, and ETH, with data from LSEG Data and Analytics. We also consider four daily traded industry portfolios: consumer goods and services, manufacturing and utilities, technology and communications, as well as healthcare, medical equipment, and drugs; these were obtained from Kenneth French's data library [Fre24]. The data spans from September 8th, 2017, to September 22nd, 2024, for a total of 2565 days, or 1729 trading days. (Although crypto assets are traded every day, we rebalance our portfolios only on market trading days; we do, however, realize gains and losses on crypto assets on weekends and holidays.) Figure 3 shows the normalized price evolution of the six assets, and table 2 lists some metrics for them. The data and code to reproduce the results are available at

https://github.com/cvxgrp/crypto_portfolio.

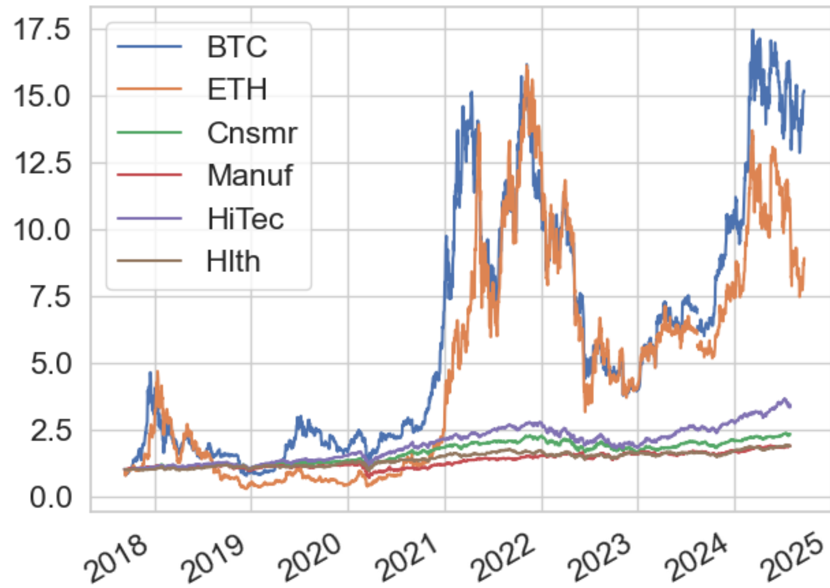

Figure 3: Normalized prices of BTC, ETH, and four industry portfolios.

Table 2: Performance metrics for the six assets.

Metric	BTC	ETH	Cnsmr	Manuf	HiTec	Hlth
Return (%)	43.5	47.1	14.1	11.4	20.7	10.8
Volatility (%)	58.1	71.6	19.3	20.5	23.9	18.0
Sharpe	0.73	0.60	0.76	0.58	0.90	0.62
Drawdown (%)	83.3	93.9	28.5	42.7	35.4	26.8

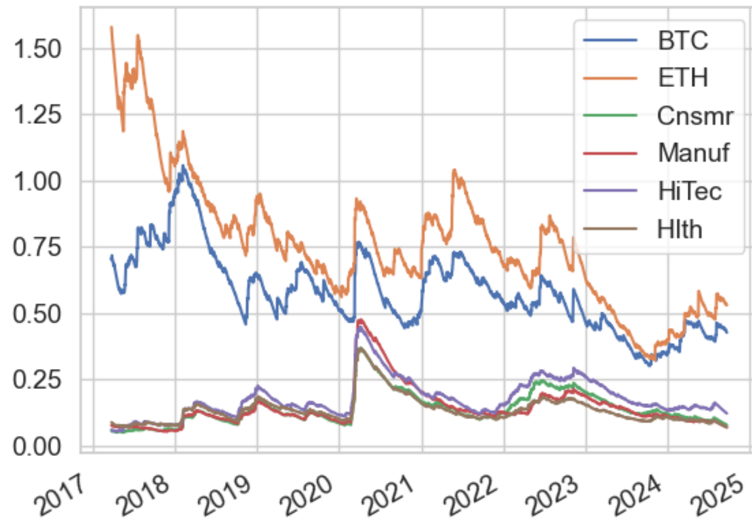

Figure 4: Estimated annualized volatilities of the six assets.

Risk model. We estimate the covariance matrix of the assets using an iterated EWMA, described in detail in [JOP+23, §2.5]. We use a 63-day half-life for the volatility estimate and a 125-day half-life for the correlation estimate. The estimated volatilities of the six assets are shown in figure 4. Two examples (from different time-periods) of the estimated correlation matrix are shown in figure 5.

Simulation and parameters. We simulate three portfolios.

- *Industries* contains only the four industry portfolios and cash.
- *Crypto* contains only the two crypto assets and cash.
- *Combined* contains all six assets and cash.

We rebalance the portfolios every trading day, using risk parity. We impose a 10% annualized risk limit on the portfolio, *i.e.*, $\sigma = 0.1\sqrt{D}$, where $D = 250$ is the number of trading days in a year. To estimate the risk of the unconstrained risk parity portfolio $x^\star$, we use a 10-day half-life EWMA of the portfolio return. We also impose a 10% maximum weight constraint on crypto assets, *i.e.*, for BTC and ETH combined. These limits were chosen as reasonable values that one might use in practice; the results are not sensitive to these choices.

4.2 Metrics

We describe the metrics used to evaluate the performance of the portfolios over the time interval $t = 1, \ldots, T$.

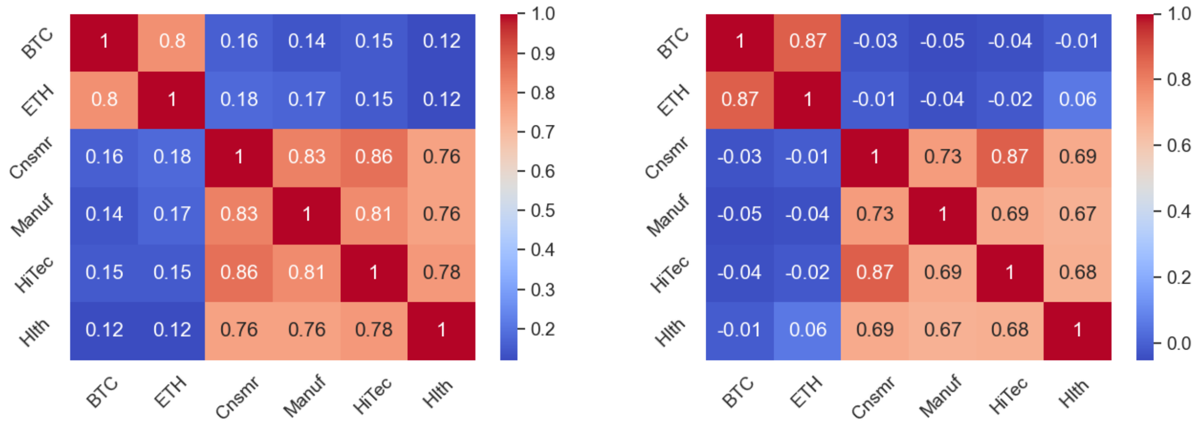

(a) March 1st, 2020. **(b)** May 1st, 2023.

Figure 5: Estimated correlation matrices on two different dates.

Return. The (realized) return of the portfolio at time t is given by

$$w_t^T r_t,$$

where r_t and w_t are the vector of (realized) asset returns and the portfolio weights at time t, respectively. The annualized (realized) return is given by

$$\bar{r} = \frac{D}{T} \sum_{t=1}^{T} w_t^T r_t.$$

Volatility. The (realized) annualized volatility of the portfolio is given by

$$\left(\frac{D}{T} \sum_{t=1}^{T} (r_t - \bar{r})^2 \right)^{1/2}.$$

Sharpe ratio. The Sharpe ratio is the ratio of the annualized return to the annualized volatility.

Drawdown. Let V_t denote the portfolio value in time period t, starting from $V_1 = 1$, with returns compounded or re-invested. These are found from the recursion $V_{t+1} = (1 + w_t^T r_t) V_t$, $t = 1, \ldots, T-1$. The (maximum) drawdown of the portfolio is defined as

$$\max_{1 \le t_1 < t_2 \le T} \left(1 - \frac{V_{t_2}}{V_{t_1}} \right),$$

i.e., the maximum fractional drop in value form a previous high.

Table 3: Portfolio performance metrics.

Metric	Industries	Crypto	Combined
Return (%)	6.0	4.5	8.2
Volatility (%)	8.2	6.0	8.2
Sharpe	0.73	0.75	1.00
Drawdown (%)	12.5	15.9	19.6

4.3 Results

Portfolio weights. The portfolio weights of the three portfolios are shown in figure 6. As expected, the crypto portfolio holds mostly cash, due to the 10% crypto limit. The industry and combined portfolios have more diversified weights, varying over time. As expected these portolios hold a lot of cash during the turbulent 2020 period. On average, the industry portfolio holds 25% cash, the crypto portfolio 90% cash, and the combined portfolio 33% cash.

Performance. Figure 7 shows the value of the three portfolios over time, normalized to one at the start of the simulation. (SR denotes Sharpe ratio.) The aggregate performance of the three portfolios is shown in table 3. Including crypto assets in the portfolio clearly gives a noticeable boost, despite allocating less than 10% of the portfolio to them. The combined portfolio has a higher return and Sharpe ratio than the industry portfolio, while having similar volatility and drawdown. To get a better understanding of the performance differences, figure 8 shows the return, volatility, and Sharpe ratio of the three portfolios over each year of the simulation.

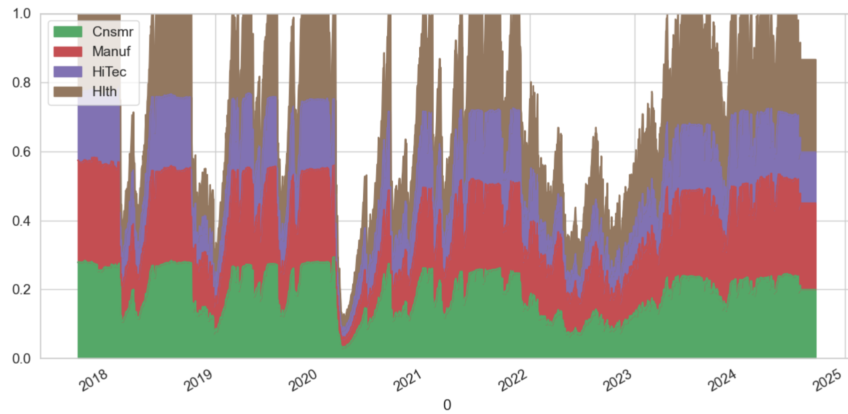

(a) Industry portfolio.

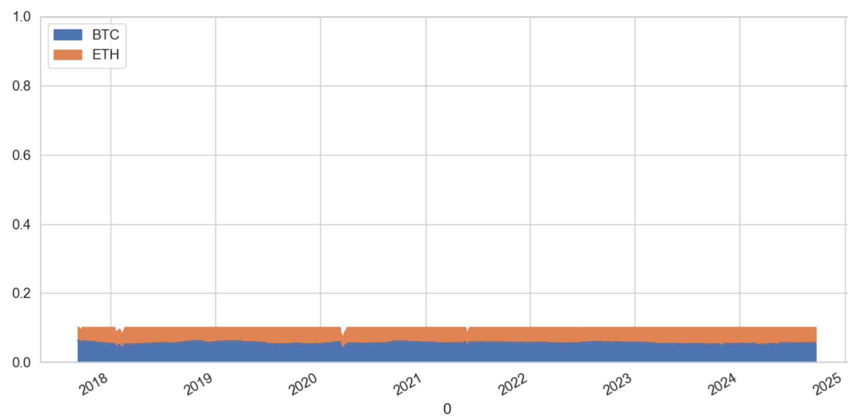

(b) Crypto portfolio.

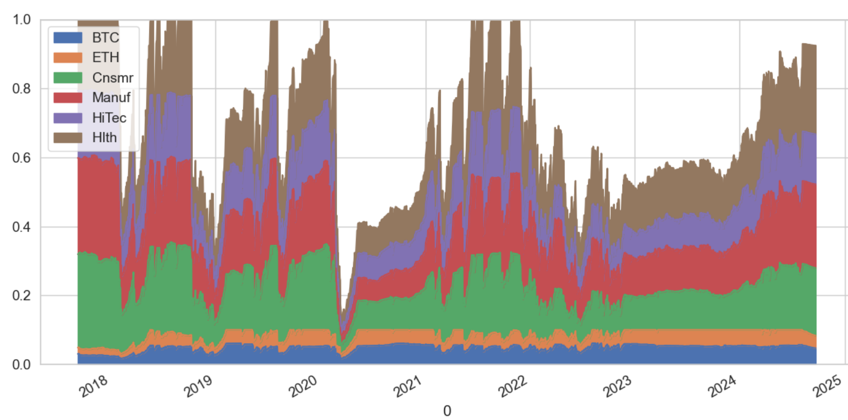

(c) Combined portfolio.

Figure 6: Portfolio weights of the three portfolios. The cash weight is shown as uncolored.

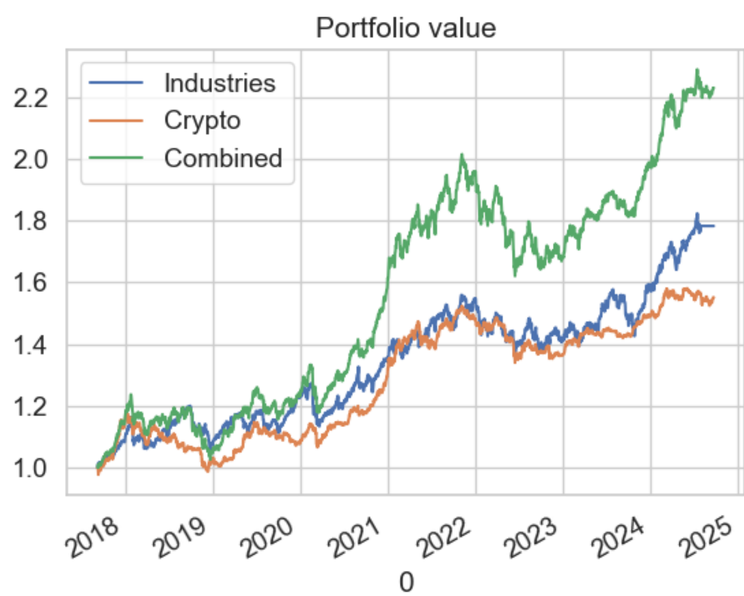

Figure 7: Portfolio values of the three portfolios.

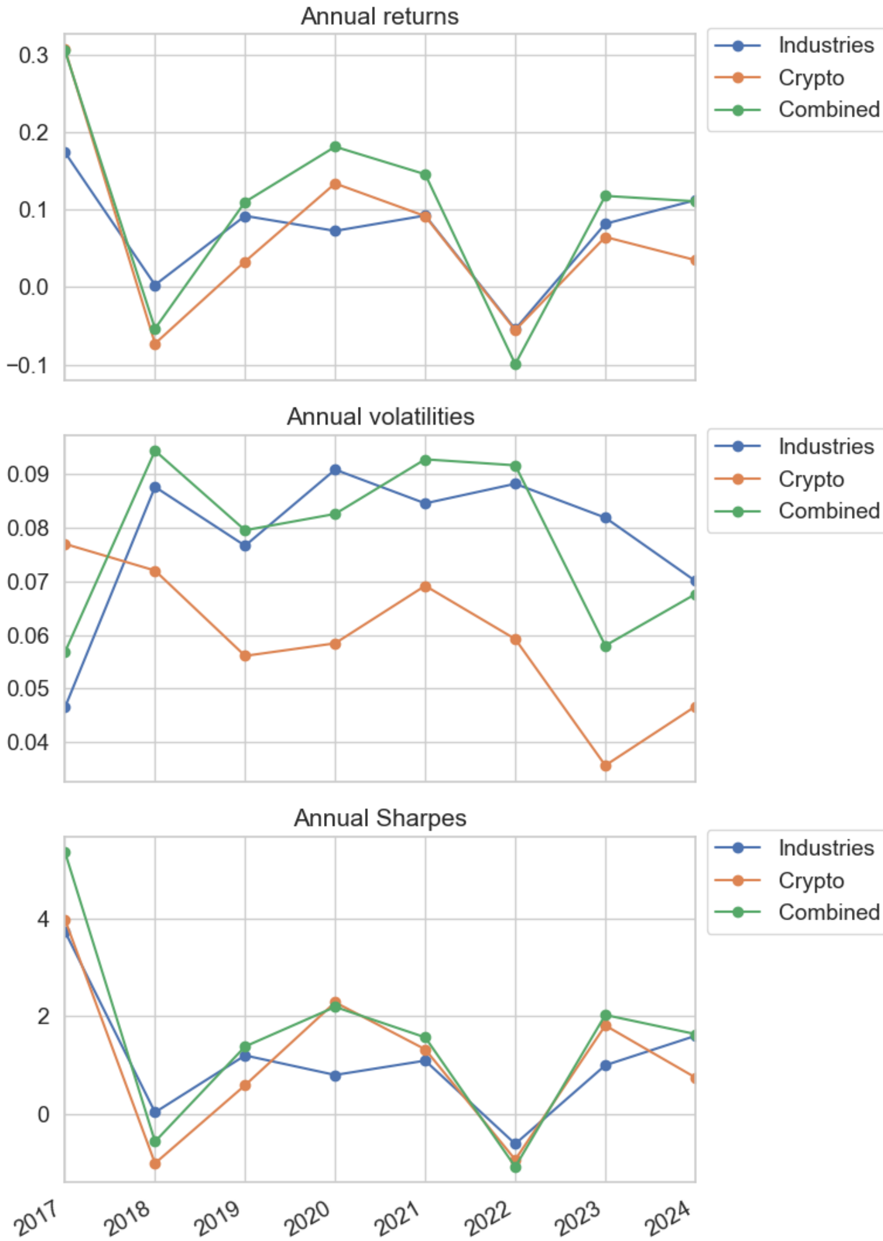

Figure 8: Annual performance metrics of the three portfolios.

Table 4: Shapley attributions by asset category.

	Cnsmr	Manuf	HiTec	Hlth	Crypto	Total
Return (%)	1.8	0.7	2.5	0.8	2.5	8.2
Volatility (%)	1.8	1.8	2.0	1.9	0.7	8.2
Sharpe	0.20	0.06	0.26	0.08	0.40	1.0
Drawdown (%)	3.3	2.4	3.2	3.5	7.3	19.6

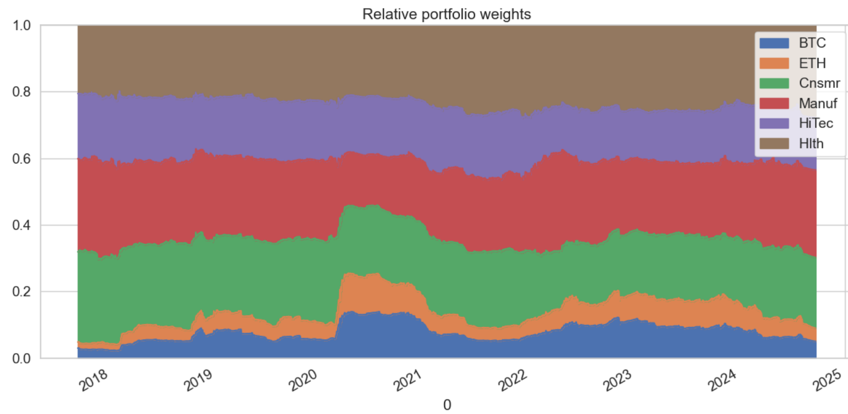

Figure 9: Relative weights of the combined portfolio.

4.4 Shapley attributions

We would like to attribute the performance of the portfolio to the different asset classes. Shapley values account for each assets's contribution to the portfolio, ensuring a fair allocation. They are uniquely characterized by satisfying a collection of desirable properties, including fairness, monotonicity, and full attribution [Sha53, HS12, ZSGJ23, FSN21, OP17, MBA21]. We will now look at the Shapley attribtutions to each industry and to crypto as a whole, for the combined portfolio. The Shapley attributions of the different asset classes for the combined portfolio are shown in table 4. Crypto assets have the highest attribution to return and Sharpe. All assets other than crypto have around a 2% contribution to volatility; crypto has a noticeably lower volatility contribution. Crypto assets have the highest contribution to drawdown.

5 Dynamically diluted 90/10 portfolio

Figure 9 shows the relative non-cash weights, *i.e.*, $w/\mathbf{1}^T w$ for the combined portfolio over time. We see that, apart from 2020, the relative weights are relatively stable and evenly distributed with about 10% in crypto assets (equally split between BTC and ETH), and 90% roughly equally split between the four industry portfolios. This motivates an even simpler portfolio construction method, akin to the popular 60/40 stocks/bonds allocation.

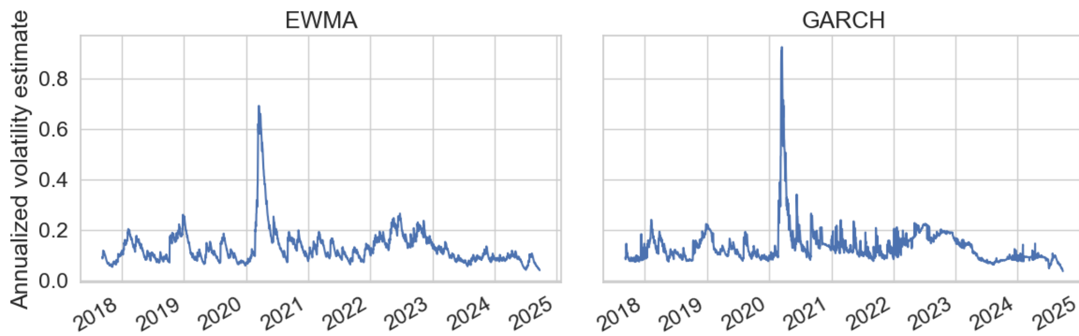

Figure 10: Annualized volatility estimates of the 90/10 portfolio.

- *90/10 portfolio.* Construct a portfolio consisting of 90% equities (*e.g.*, the four industries with equal weights) and 10% crypto (*e.g.*, equally split between BTC and ETH).
- *Dynamic cash dilution.* Based on an estimate of the recent volatility of the 90/10 portfolio, dilute the 90/10 portfolio with cash to achieve the target risk σ and respect weight limits. We refer to this portfolio as the dynamically diluted 90/10 (DD90/10) portfolio.

5.1 DD90/10 results

Volatility estimators. We evaluate the performance of the DD90/10 portfolio using two different volatility estimators: a 10-day half-life EWMA, and a GARCH(1,1) model refitted every day on the last 250 days of data, using the arch package in Python [She24]. (We tried several other volatility estimators; all gave similar results.) The volatility estimates of the 90/10 portoflio are shown in figure 10.

Weights. The weights of the DD90/10 portfolios are shown in figure 11. The weights are quite similar, with the GARCH estimator being noticeably more reactive during some periods.

Performance. The performance of the DD90/10, compared to the (combined) CRA portfolio, is shown in table 5. Figure 12 shows the value of the three portfolios over time.

6 Conclusions

We have illustrated that, despite the documented extreme behaviors of crypto assets, simple traditional portfolio contrstruction techniques can be used to include them in a diversified portfolio. We show this using two standard portfolio construction methods, one based on

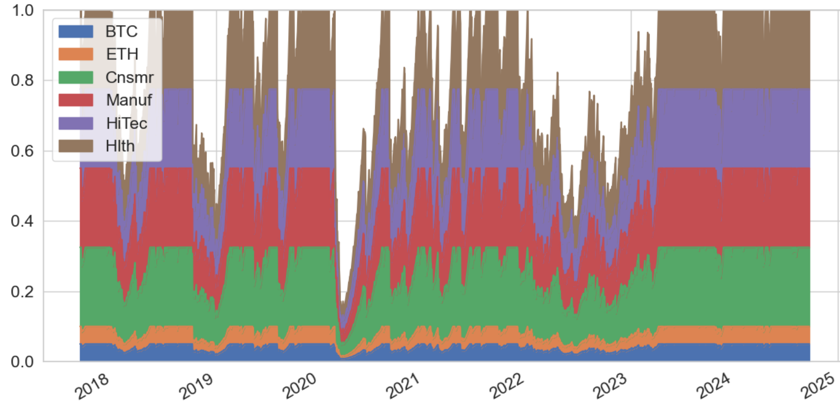

(a) EWMA volatility estimator.

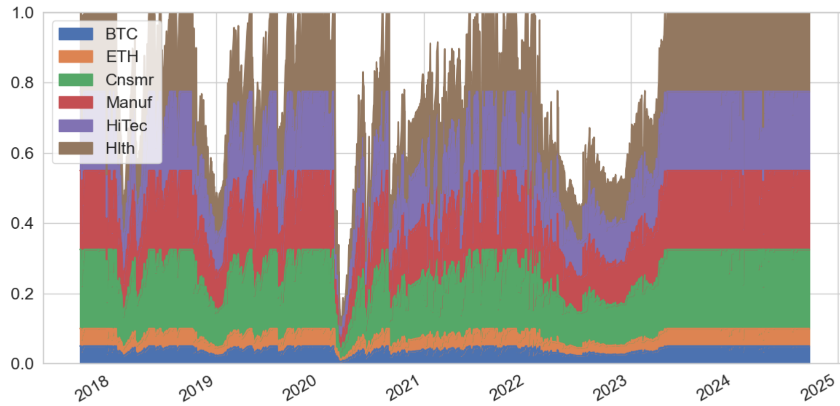

(b) GARCH volatility estimator.

Figure 11: Portfolio weights of the DD90/10 portfolios with EWMA and GARCH volatility estimators.

Metric	DD90/10 (EWMA)	DD90/10 (GARCH)	CRA
Return (%)	10.4	10.1	8.2
Volatility (%)	9.8	9.7	8.2
Sharpe	1.06	1.04	1.00
Drawdown (%)	19.9	19.7	19.6

Table 5: Performance metrics DD90/10 and CRA.

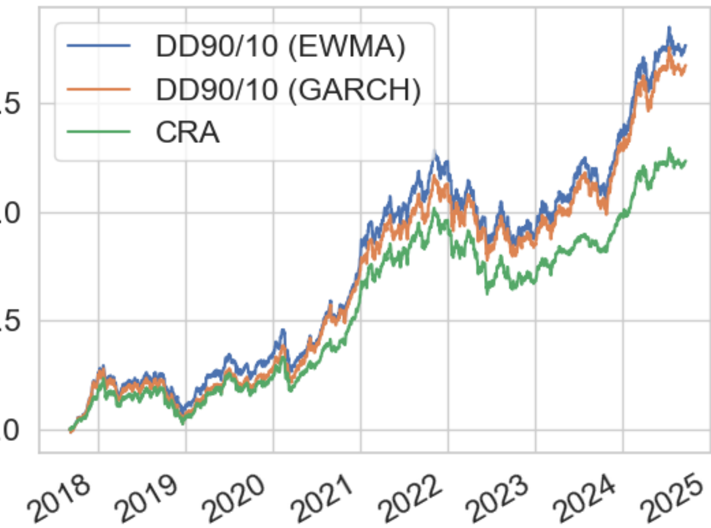

Figure 12: Portfolio values of the DD90/10 and CRA portfolios.

risk parity, and the other a fixed set of relative weights, with each one dynamically diluted with cash to achieve a target ex-ante risk. The addition of even a modest crypto weight of 10% increases the return and Sharpe ratio of the portfolio significantly, without significantly increasing volatility or drawdown.

Acknowledgments. This work was supported by the IOG Research Hub.

References

[BB06] J. Bulla and I. Bulla. Stylized facts of financial time series and hidden semi-Markov models. *Computational statistics & data analysis*, 51(4):2192–2209, 2006.

[BB22] S. Barratt and S. Boyd. Covariance prediction via convex optimization. *Optimization and Engineering*, 2022.

[BBDRM21] S. Bartram, J. Branke, G. De Rossi, and M. Motahari. Machine learning for active portfolio management. *Journal of Financial Data Science*, 3(3):9–30, 2021.

[BJK+24] S. Boyd, K. Johansson, R. Kahn, P. Schiele, and T. Schmelzer. Markowitz portfolio construction at seventy. *Journal of Portfolio Management*, 50(8):117–160, 2024.

[BL90] F. Black and R. Litterman. Asset allocation: Combining investor views with market equilibrium. *Goldman Sachs Fixed Income Research*, 115(1):7–18, 1990.

[BL92] F. Black and R. Litterman. Global portfolio optimization. *Financial analysts journal*, 48(5):28–43, 1992.

[Bla89] F. Black. Universal hedging: Optimizing currency risk and reward in international equity portfolios. *Financial Analysts Journal*, 45(4):16–22, 1989.

[BM19] A. Brauneis and R. Mestel. Cryptocurrency-portfolios in a mean-variance framework. *Finance Research Letters*, 28:259–264, 2019.

[Bol86] T. Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327, 1986.

[BRAMEK21] W. Bakry, A. Rashid, S. Al-Mohamad, and N. El-Kanj. Bitcoin and portfolio diversification: A portfolio optimization approach. *Journal of Risk and Financial Management*, 14(7):282, 2021.

[Bur19] T. Burggraf. Risk-based portfolio optimization in the cryptocurrency world. *Available at SSRN 3454764*, 2019.

[BV04] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[BV24] S. Boyd and L. Vandenberghe. Convex optimization additional exercises. https://github.com/cvxgrp/cvxbook_additional_exercises, 2024.

[CC08] Y. Choueifaty and Y. Coignard. Toward maximum diversification. *The Journal of Portfolio Management*, 35(1):40–51, 2008.

[CK13] J. Ciolli and J. Kara. Paul Krugman: Bitcoin is a more obvious bubble than housing was. *Business Insider*, 2013.

[CK23] S. Chakrabarty and A. Kanaujiya. *Mathematical Portfolio Theory and Analysis*. Springer Nature, 2023.

[CT06] G. Cornuejols and R. Tütüncü. *Optimization methods in finance*, volume 5. Cambridge University Press, 2006.

[DB16] S. Diamond and S. Boyd. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.

[EB86] R. Engle and T. Bollerslev. Modelling the persistence of conditional variances. *Econometric reviews*, 5(1):1–50, 1986.

[Eng82] R. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica: Journal of the econometric society*, pages 987–1007, 1982.

[Eng02] R. Engle. Dynamic conditional correlation. *Journal of Business & Economic Statistics*, 20(3):339–350, 2002.

[FP15] Y. Feng and D. Palomar. Scrip: Successive convex optimization methods for risk parity portfolio design. *IEEE Transactions on Signal Processing*, 63(19):5285–5300, 2015.

[Fre24] K. French. Kenneth french data library. https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html#Research, 2024.

[FSN21] D. Fryer, I. Strümke, and H. Nguyen. Shapley values for feature selection: The good, the bad, and the axioms. *IEEE Access*, 9:144352–144360, 2021.

[GB23] A. Gunjan and S. Bhattacharyya. A brief review of portfolio optimization techniques. *Artificial Intelligence Review*, 56(5):3847–3886, 2023.

[GBWZ23] B. Ghosh, E. Bouri, J. Bum Wee, and N. Zulfiqar. Return and volatility properties: Stylized facts from the universe of cryptocurrencies and NFTs. *Research in International Business and Finance*, 65:101945, 2023.

[Gea36] R. Geary. Moments of the ratio of the mean deviation to the standard deviation for normal samples. *Biometrika*, pages 295–307, 1936.

[Gea47] R. Geary. Testing for normality. *Biometrika*, 34(3/4):209–242, 1947.

[GK00] R. Grinold and R. Kahn. *Active portfolio management*. McGraw-Hill, 2000.

[GPP20] P. Giudici, P. Pagnottoni, and G. Polinesi. Network models to enhance automated cryptocurrency portfolio management. *Frontiers in artificial intelligence*, 3:22, 2020.

[Gur22] I. Gurrib. Machine learning and portfolio management: A review. *Annals of Mathematics and Computer Science*, 5:31–43, 2022.

[Hol20] O. Holovatiuk. Cryptocurrencies as an asset class in portfolio optimisation. *Central European Economic Journal*, 7(54):33–55, 2020.

[HPR19] A. Hu, C. Parlour, and U. Rajan. Cryptocurrencies: Stylized facts on a new investible instrument. *Financial Management*, 48(4):1049–1068, 2019.

[HRF19] Y. Hu, S. Rachev, and F. Fabozzi. Modelling crypto asset price dynamics, optimal crypto portfolio, and crypto option valuation. *arXiv preprint arXiv:1908.05419*, 2019.

[HS12] F. Huettner and M. Sunder. Axiomatic arguments for decomposing goodness of fit according to Shapley and Owen values. 2012.

[JL17] Z. Jiang and J. Liang. Cryptocurrency portfolio management with deep reinforcement learning. In *2017 Intelligent systems conference (IntelliSys)*, pages 905–913. IEEE, 2017.

[JOP+23] K. Johansson, G. Ogut, M. Pelger, T. Schmelzer, and S. Boyd. A simple method for predicting covariance matrices of financial returns. *Foundations and Trends in Econometrics*, 12(4):324–407, 2023.

[JWZZ20] L. Jiang, K. Wu, G. Zhou, and Y. Zhu. Stock return asymmetry: Beyond skewness. *Journal of Financial and Quantitative Analysis*, 55(2):357–386, 2020.

[JXL17] Z. Jiang, D. Xu, and J. Liang. A deep reinforcement learning framework for the financial portfolio management problem. *arXiv preprint arXiv:1706.10059*, 2017.

[KRNY24] M. Karim, Mohamed E. Rasid, A. Noman, and L. Yarovaya. Exploring asymmetries in cryptocurrency intraday returns and implied volatility: New evidence for high-frequency traders. *International Review of Financial Analysis*, page 103617, 2024.

[Kru11] P. Krugman. Golden cyberfetters. *The New York Times*, 2011.

[Kru13] P. Krugman. Adam Smith hates bitcoin. *Business Insider*, 2013.

[KTF14] P. Kolm, R. Tütüncü, and F. Fabozzi. 60 years of portfolio optimization: Practical challenges and current trends. *European Journal of Operational Research*, 234(2):356–371, 2014.

[KX23] B. Kelly and D. Xiu. Financial machine learning. *Foundations and Trends® in Finance*, 13(3-4):205–363, 2023.

[LB20] G. Lucarelli and M. Borrotti. A deep Q-learning portfolio management framework for the cryptocurrency market. *Neural Computing and Applications*, 32:17229–17244, 2020.

[LT21] Y. Liu and A. Tsyvinski. Risks and returns of cryptocurrency. *The Review of Financial Studies*, 34(6):2689–2727, 2021.

[Mar52] H. Markowitz. Portfolio selection. *Journal of Finance*, 7(1):77–91, 1952.

[MBA21] N. Moehle, S. Boyd, and A. Ang. Portfolio performance attribution via shapley value. *arXiv preprint arXiv:2102.05799*, 2021.

[Meu10] A. Meucci. Fully flexible views: Theory and practice. *arXiv preprint arXiv:1012.2848*, 2010.

[Mic89] R. Michaud. The Markowitz optimization enigma: Is 'optimized' optimal? *Financial analysts journal*, 45(1):31–42, 1989.

[MT10] H. Malmsten and T. Teräsvirta. Stylized facts of financial time series and three popular models of volatility. *European Journal of pure and applied mathematics*, 3(3):443–477, 2010.

[Nak08] S. Nakamoto. A peer-to-peer electronic cash system. *Bitcoin.–URL: https://bitcoin. org/bitcoin. pdf*, 4(2):15, 2008.

[Nar24] R. Narang. *Inside the Black Box: A Simple Guide to Systematic Investing.* J. Wiley & Sons, 2024.

[OP17] A. Owen and C. Prieur. On Shapley value for measuring importance of dependent inputs. *SIAM/ASA Journal on Uncertainty Quantification*, 5(1):986–1002, 2017.

[Pal24] D. Palomar. *Portfolio Optimization.* Cambridge University Press, 2024.

[PU19] E. Platanakis and A. Urquhart. Portfolio management with cryptocurrencies: The role of estimation risk. *Economics Letters*, 177:76–80, 2019.

[Qia11] E. Qian. Risk parity and diversification. *The Journal of Investing*, 20(1):119–127, 2011.

[Ram21] G. Ramkumar. Cryptocurrency portfolio construction using machine learning models. In *Contemporary Trends and Challenges in Finance: Proceedings from the 6th Wroclaw International Conference in Finance*, pages 103–122. Springer, 2021.

[RRdSJB22] I. Ruiz Roque da Silva, E. Junior, and P. Balbi. Cryptocurrencies trading algorithms: A review. *Journal of Forecasting*, 41(8):1661–1668, 2022.

[RU00] T. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–42, 2000.

[Sha53] L. Shapley. A value for n-person games. *Contribution to the Theory of Games*, 2, 1953.

[She24] K. Sheppard. Arch documentation. 2024.

[WFP20] L. Wu, Y. Feng, and D. Palomar. General sparse risk parity portfolio design via successive convex optimization. *Signal Processing*, 170:107433, 2020.

[Yar13] J. Yarow. Paul Krugman responds to all the people throwing around his old internet quote. *Business Insider*, 2013.

[Yer24] D. Yermack. Is bitcoin a real currency? An economic appraisal. In *Handbook of digital currency*, pages 29–40. Elsevier, 2024.

[ZSGJ23] H. Zhang, H. Singh, M. Ghassemi, and S. Joshi. Why did the model fail? Attributing model performance changes to distribution shifts. 2023.