Università
degli Studi
di Padova

University of Padua
Master Thesis

Portfolio optimization with dominant cryptocurrencies: an empirical study of risk, return, and diversification benefits

Francesco Landi
ID 2103413

Supervisor
Prof. Massimiliano Caporin
Department of Statistical Sciences, University of Padova

Academic Year 2024/2025

Contents

Abstract

This study investigates the role of dominant cryptocurrencies, selected based on relative market-cap thresholds of 1% and 5%, in enhancing portfolio performance, risk management, and diversification. Using daily data from 2015 to 2023 sourced from CoinGecko and Refinitiv, the analysis evaluates portfolio strategies that combine dominant cryptocurrencies with traditional financial assets, including equities, bonds, commodities, and gold, following the reference framework of Petukhina et al. (2021). The methodology compares passive benchmarks, mean-variance optimization, utility-based allocation, and Smart Beta 2.0 risk-based rules across crypto-only, traditional-only, and mixed-asset universes. It then tests the robustness of the frictionless results under an order-driven implementation with transaction costs and slippage. Findings show that dominant cryptocurrencies can improve portfolio outcomes only when allocation rules control estimation error, turnover, and regime dependence. The results provide practical evidence for investors seeking to balance digital-asset innovation with portfolio stability.

1 Introduction

1.1 Motivation and context

Cryptocurrencies started as a niche experiment in peer-to-peer digital money. That framing no longer holds. CoinGecko's 2023 industry review states the rebound from 2022 in round numbers: over the calendar year, total cryptocurrency market capitalization rose by about 108%, from roughly USD 829 billion to roughly USD 1.72 trillion (CoinGecko 2024). The 2025 instalment of the same report line reads less like a steady climb and more like a whip: CoinGecko flags a short aggregate cap above USD 4 trillion in late 2025, then a sharp fourth-quarter sell-off that leaves year-end headline totals nearer USD 3 trillion, with the full-year change negative on their reckoning (CoinGecko 2026). The empirical work here stops in 2023. Those headline totals therefore refer to the same year as the final observations in the sample, not to a later run-up in size. They still do not put cryptocurrencies on equal footing with equities, bonds, or commodities. But they do make the asset class too large to treat as a peripheral curiosity. The real question for investors is whether its return and risk profile can be built into disciplined allocation rules at all.

The natural starting point is modern portfolio theory. The mean-variance framework of Markowitz (1952) formalizes the trade-off between expected return and variance, and it remains the reference language for empirical asset allocation. Applying it to cryptocurrencies is not automatic, though. Digital assets trade continuously, experience abrupt regime shifts, and produce fat-tailed returns that undermine the smooth assumptions on which many classical portfolio intuitions rest. The empirical task of this thesis is therefore not to assume that standard allocation methods work in crypto markets, but to test which parts of the portfolio-theoretic toolkit survive when the investable universe is dynamic, volatile, and operationally costly to rebalance.

The market structure around cryptocurrencies also changed materially in 2024. On January 10, 2024, the U.S. Securities and Exchange Commission approved the listing and trading of a number of spot bitcoin exchange-traded product shares (Gensler 2024). This approval did not remove the risks of the underlying asset. The SEC framed the products through disclosure and exchange-listing standards, not as an endorsement of bitcoin itself. Still, it created a regulated, exchange-traded channel through which traditional investors could obtain spot bitcoin exposure. That institutional access point makes the comparison between cryptocurrencies and conventional ETF exposures more practically relevant than it was before.

The central empirical challenge is that the diversification role of cryptocurrencies is unstable. Volatility is high. Correlations with traditional assets shift across regimes. The same asset can act as a diversifier in one period and as a source of concentrated tail risk in another. This thesis studies that instability directly by combining a dynamic cryptocurrency universe with a static traditional-asset benchmark, then evaluating both frictionless and cost-inclusive implementations.

1.2 Research gap

The literature on cryptocurrency allocation has shown that digital assets can alter the risk-return profile of diversified portfolios, but many empirical designs rely on fixed lists of large coins or on sample windows that do not cover the full sequence of recent market regimes. Petukhina et al. (2021) provide a natural baseline for studying cryptocurrency portfolio allocation through systematic rules. This thesis extends that logic in three directions.

First, the sample covers 2015–2023, spanning the early expansion of tradable crypto markets, the 2017 bull market, the 2018 contraction, the 2021 speculative surge, the 2022 crash, and the 2023 recovery. Second, the cryptocurrency sleeve is not a fixed top-N list. It is selected dynamically using a rolling 12-month activity requirement and a 1% relative market-capitalization threshold within the selected macro-category, as documented in section 3.5. Third, the analysis does not stop at frictionless portfolio accounting. Chapter 6 adds an order-driven implementation layer with OHLC paths, transaction costs, and slippage assumptions.

The conceptual gap matters too. Return-based optimizers, especially the Maximum Sharpe Ratio portfolio, depend heavily on sample mean estimates. This is precisely the setting in which mean-variance optimization can become an "error maximizer" (Michaud 1989). Chopra and Ziemba (1993) further show that errors in expected returns are especially damaging for optimal portfolio choice. Cryptocurrency returns make this problem worse because sample means are unstable across regimes and extreme observations can dominate rolling windows. The empirical results in section 5.2 and Appendix B use this theory to interpret the fragility of return-seeking allocations.

1.3 Cryptocurrencies as an asset class

Cryptocurrencies are ledger-based digital assets whose ownership is recorded on distributed networks rather than on the books of a centralized issuer or custodian. Unlike equities or bonds, most cryptocurrencies do not grant a contractual claim on cash flows. Valuation is therefore driven by network use, monetary design, liquidity, speculative demand, and the perceived usefulness of the protocol or ecosystem. That makes them difficult to place within the standard asset-class taxonomy, even when they are traded through increasingly institutional market infrastructure.

The empirical signature of the asset class is high volatility and strong non-normality. Chapter 3 documents this directly through descriptive statistics, density plots, and diagnostic figures for the selected cryptocurrency universe. These features are not cosmetic details. They affect every portfolio statistic used later in the thesis, from covariance estimates to adjusted Sharpe ratios and certainty-equivalent returns. A strategy that performs well under Gaussian assumptions may perform poorly when returns are dominated by jumps, crash-recovery cycles, and time-varying liquidity.

Liquidity has improved substantially since the early years of cryptocurrency trading, but it remains uneven across assets and regimes. Large coins such as Bitcoin and Ethereum are supported by deep centralized and decentralized markets. Smaller tokens can become illiquid or economically irrelevant after a short period of attention. This motivates the dynamic universe design used here. Rather than assuming that today's large assets were investable throughout the full sample, the selection rule admits assets only when they satisfy backward-looking activity and market-capitalization criteria.

Correlations with traditional assets must also be read as regime-dependent. Prior work on bitcoin diversification finds that digital assets can contribute to portfolio diversification, but that the benefit is sensitive to the allocation rule and sample period (Brière, Oosterlinck, and Szafarz 2015). The mixed-asset analysis in Chapter 5 follows this logic by allowing the optimizer to choose between crypto and traditional ETF exposures at each rebalance, rather than imposing a fixed crypto sleeve.

1.4 Dynamic versus static universe design

A static top-N cryptocurrency universe is easy to reproduce, but it can embed substantial look-ahead and survivorship bias: continuity over the look-back window is required, so listings and delistings are handled at each rebalance date instead of being written back into earlier dates. Assets that are large at the end of the sample may have been unavailable, illiquid, or economically marginal at earlier dates. Assets that mattered historically may disappear from a current top-N list. A static universe therefore risks measuring the performance of a retrospectively convenient basket rather than an investable portfolio process.

The rolling-threshold design used in this thesis addresses this by making eligibility time-dependent, so the investable set changes month by month with the filters rather than staying on one fixed list for the entire sample. At each rebalance date T, the selection procedure uses only information available up to $T-1$. A cryptocurrency must remain continuously active over the preceding 365 days and must exceed the relative market-capitalization threshold[1] within the retained macro-category. This rule allows the universe to adapt to structural change while preserving temporal discipline. Examples such as Polkadot, Solana, and layer-2 ecosystems[2] show why adaptability matters: innovation can enter the opportunity set, but only after it becomes sufficiently liquid and observable.

The design also excludes assets whose mechanics are inconsistent with the allocation problem studied here. Stablecoins are excluded because their purpose is price anchoring rather than risky capital appreciation. Wrapped and bridged

[1] A relative market-capitalization threshold requires an asset to exceed a fixed share of its macro-category's total market capitalization, which keeps the selected assets liquid and economically meaningful.

[2] Layer-2 ecosystems are scalability solutions built atop layer-1 blockchains, for example Arbitrum or Optimism for Ethereum, that inherit base-layer security while reducing transaction costs through off-chain or batched processing.

assets[3] are excluded to avoid duplicate economic exposures. Meme tokens and other categories with weak investment fundamentals are removed by the macro-category filter described in Chapter 3. The traditional universe follows the opposite logic: the seven ETFs in Table 3.4 are held fixed because they represent mature and persistent exposures to equities, bonds, commodities, and gold.

1.5 Portfolio strategies overview

The strategy set combines passive benchmarks, classical mean-variance optimization, and Smart Beta 2.0 risk-based rules. The equally weighted portfolio is the simplest benchmark: it assigns the same capital weight to every eligible asset and avoids estimation risk. The value-weighted portfolio assigns weights according to market capitalization and represents the scale-concentrated benchmark commonly used in market-index construction.

The Markowitz family begins with the Global Minimum Variance portfolio, which uses only the covariance matrix and selects the lowest-variance long-only allocation. The Maximum Sharpe Ratio portfolio adds expected returns and seeks the best excess-return-to-risk trade-off, making it more sensitive to estimation error. Utility-maximizing portfolios generalize this trade-off by introducing a risk-aversion parameter γ, so that more risk-tolerant investors can accept higher variance in exchange for higher expected return.

The Smart Beta 2.0 family (Amenc, Goltz, and Martellini 2013) shifts attention from expected returns to diversification and risk structure. Maximum Decorrelation minimizes average portfolio correlation. Maximum Diversification Ratio maximizes the gap between weighted individual volatilities and total portfolio volatility. Equal Risk Contribution allocates capital so that each asset contributes equally to portfolio risk. Diversified Risk Parity is a simpler inverse-volatility rule derived under an equal-correlation approximation. Formal definitions, constraints, and evaluation metrics appear in Chapter 4.

1.6 Why crypto breaks naive mean-variance

The core weakness of naive mean-variance allocation in cryptocurrency markets is the instability of expected returns. Rolling sample means can change sign across short windows, and isolated crashes or rallies can dominate the estimate

[3] Wrapped tokens map a native asset onto another chain via smart contracts (for example wrapped Bitcoin on Ethereum) and add protocol and counterparty risk relative to holding the native coin. Bridged tokens move value across chains through protocols that lock, mint, or release claims; when those mechanisms fail, losses can be large.

used at the next rebalance. When the optimizer is allowed to maximize the Sharpe ratio, these noisy estimates feed directly into portfolio weights. The result is usually not a diversified expression of expected performance. It is an amplified expression of sampling error (Chopra and Ziemba 1993; Michaud 1989).

Covariance estimates are also unstable. Cryptocurrency volatility clusters in episodes and correlations can tighten sharply during market stress. A rolling-window optimizer must therefore estimate both expected returns and the covariance matrix in precisely the periods when the data-generating process is changing most. This is why the thesis pays particular attention to risk-based strategies that reduce or eliminate reliance on the expected-return vector.

The empirical consequences are visible in Chapter 5. In the crypto-only scenario, return-seeking strategies suffer from unstable drawdowns and weak adjusted performance, while risk-based rules remain more tractable. The case study in Appendix B is the clearest example of this. A single regime shift can turn an apparently attractive return-based allocation into concentrated collapse risk.

2 Background

2.1 Mean–variance portfolio theory and its limitations

Modern portfolio theory begins with the mean–variance framework introduced by Markowitz (Markowitz 1952). For a universe of N risky assets, let $\boldsymbol{\mu} \in \mathbb{R}^N$ denote expected excess returns, $\boldsymbol{\Sigma} \in \mathbb{R}^{N\times N}$ the covariance matrix, and $\mathbf{w} \in \mathbb{R}^N$ the vector of portfolio weights. The canonical long-only mean–variance problem can be written as

$$\max_{\mathbf{w}} \quad \mathbf{w}'\boldsymbol{\mu} - \frac{\gamma}{2}\mathbf{w}'\boldsymbol{\Sigma}\mathbf{w} \quad \text{s.t.} \quad \mathbf{1}'\mathbf{w} = 1, \quad \mathbf{w} \geq 0, \tag{2.1}$$

where $\gamma > 0$ is the investor's risk-aversion parameter. Equivalent formulations minimize variance for a target return, maximize return for a target variance, or solve for the tangency portfolio. The set of non-dominated portfolios forms the efficient frontier: for a given expected return, no portfolio on the frontier has lower variance, and for a given variance, no portfolio has higher expected return (Markowitz 1952).

The practical weakness of this framework is that the inputs are not observed. They must be estimated from finite samples, and optimization converts input error into allocation error. This is the core of the error-maximization critique: an optimizer tends to overweight assets with upward-biased estimated returns, downward-biased estimated variances, or unusually favorable estimated correlations, so the most attractive solution in sample may be precisely the least reliable out of sample (Michaud 1989). The problem is structural rather than computational. Even if the quadratic program is solved exactly, the solution inherits the uncertainty of $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\Sigma}}$.

Estimation error is a first-order concern because mean–variance weights are highly sensitive to the expected-return vector. In the unconstrained case, optimal weights are proportional to $\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}$, so small changes in $\boldsymbol{\mu}$ can produce large reallocations when assets are correlated or when the covariance matrix is ill-conditioned. Errors in expected returns are typically more damaging than errors in variances and covariances, though all three inputs matter for portfolio choice (Chopra and Ziemba 1993). This distinction explains why purely return-sensitive portfolios, such as tangency and utility portfolios, often require stronger assumptions than risk-based allocations.

These limitations are especially relevant in empirical cryptocurrency allocation. Short histories, extreme observations, regime shifts, and changing market composition make the estimated moments unstable. The theoretical lesson for the rest of the thesis is not that mean–variance theory should be discarded, but that its empirical use must be benchmarked against allocation rules that reduce dependence on noisy first-moment estimates and against diagnostics that reveal out-of-sample fragility.

2.2 Smart Beta 2.0 and risk-based allocation

Smart Beta 2.0 describes a family of allocation approaches designed to separate portfolio construction from market-capitalization weighting and from unstable expected-return forecasts. In this literature, portfolio weights are built to capture diversification or risk premia through transparent, rules-based objectives rather than through discretionary return prediction (Amenc, Goltz, and Martellini 2013). The motivation connects directly to the weaknesses of mean–variance optimization: if expected returns are difficult to estimate reliably, strategies built primarily on covariance, volatility, correlation, or risk contribution may hold up better.

The Global Minimum Variance (GMV) portfolio is the most direct risk-based descendant of the Markowitz framework. It solves

$$\min_{\mathbf{w}} \quad \mathbf{w}'\boldsymbol{\Sigma}\mathbf{w} \quad \text{s.t.} \quad \mathbf{1}'\mathbf{w} = 1, \quad \mathbf{w} \geq 0, \tag{2.2}$$

and therefore ignores expected returns entirely. Its theoretical appeal is parsimony: only the covariance matrix is required. Its drawback is that it may concentrate in low-volatility assets and can still be sensitive to covariance estimation. Maximum Decorrelation (MDC) follows a related but distinct intuition. Rather than minimizing total variance directly, it emphasizes low average correlation across holdings. This can be useful when diversification failure is driven less by individual volatility and more by assets becoming jointly exposed to the same market factor.

The Maximum Diversification Ratio (MDR) portfolio formalizes diversification as the ratio between weighted average individual volatility and total portfolio volatility:

$$DR(\mathbf{w}) = \frac{\mathbf{w}'\boldsymbol{\sigma}}{\sqrt{\mathbf{w}'\boldsymbol{\Sigma}\mathbf{w}}}, \tag{2.3}$$

where $\boldsymbol{\sigma}$ is the vector of asset volatilities. Maximizing this ratio selects portfolios that get the greatest risk reduction from combining assets (Choueifaty and Coignard 2008). MDR therefore differs from GMV because it does not simply seek the lowest volatility portfolio; it seeks the strongest diversification effect relative to the risks of the constituents.

Equal Risk Contribution (ERC) and Diversified Risk Parity (DRP) belong to the risk-parity family. ERC requires each asset to contribute the same share of portfolio variance, so concentration is measured in risk terms rather than in capital weights (Maillard, Roncalli, and Teïletche 2010). DRP is a simplified version that relies mainly on inverse-volatility weighting under stronger correlation assumptions (Amenc, Goltz, and Martellini 2013). Both approaches can be read as attempts to avoid hidden concentration in a small set of dominant risk sources. Finally, the equally weighted $1/N$ portfolio remains a hard benchmark because it is simple, transparent, and difficult for optimized strategies to beat out of sample once estimation error and turnover are considered (DeMiguel, Garlappi, and Uppal 2009). The empirical relevance of Smart Beta 2.0 is therefore comparative: each rule embodies a different answer to the same problem of how to diversify when return forecasts are unreliable.

2.3 Cryptocurrencies in portfolio research: a literature review

The early portfolio literature on cryptocurrencies focused mainly on Bitcoin as a potential diversifier for traditional assets. Bitcoin could improve the risk–return profile of a diversified portfolio over the sample studied by Briere, Oosterlinck and Szafarz, but only with small allocations because volatility and downside risk are large (Brière, Oosterlinck, and Szafarz 2015). This result established an important theme: low average correlation with traditional assets may create diversification benefits, but those benefits are conditional on position size, sample period, and risk measurement. Subsequent work broadened the question from Bitcoin alone to cryptocurrency risk factors and cross-sectional behavior. Cryptocurrency returns are not well explained by standard equity, currency, or commodity factors, and instead are related to crypto-specific demand and attention variables (Liu and Tsyvinski 2021a). This supports the view that cryptocurrencies are not simply high-beta versions of existing asset classes.

Portfolio allocation studies then moved from single-asset inclusion tests toward multi-cryptocurrency universes. The central reference for this thesis evaluates portfolio allocation strategies within a cryptocurrency universe, including naive, mean–variance, and risk-based approaches (Petukhina et al. 2021). Its contribution is to treat cryptocurrencies as a portfolio problem rather than as a single speculative asset. At the same time, this setting highlights unresolved empirical-design issues: how the investment universe is selected, whether failed or delisted assets are present, and how strategy performance changes when the market composition evolves.

Selection and survivorship bias are particularly important in cryptocurrency research. The market has experienced high token attrition, exchange delistings, relabeling, and abrupt loss of liquidity. If an empirical sample contains only assets that survived until the end of the period, historical performance can be overstated. Survivorship and delisting bias can materially distort cryptocurrency portfolio performance, especially for equal-weighted portfolios (Ammann et al. 2022). This literature motivates careful distinction between ex-ante investability and ex-post availability.

Results also vary across market cycles. The 2017 boom, the 2018 crash, the 2021 expansion, and the 2021–2022 drawdown produced very different combinations of momentum, correlation, liquidity, and investor participation. Studies focusing on one cycle may reach conclusions that do not generalize to another. This explains why the literature contains both optimistic findings about diversification and more cautious evidence about instability, tail dependence, and implementation. The open question is not whether cryptocurrencies can ever improve portfolio outcomes, but under what universe definition, allocation rule, and market regime such benefits survive realistic empirical scrutiny.

2.4 The Terra/LUNA collapse and tail risk in crypto

Cryptocurrency portfolio theory cannot rely only on variance because extreme events are central to the asset class. Tail risk refers to losses in the far left tail of the return distribution, where ordinary volatility measures may give an incomplete description of investor exposure. Empirical studies document that cryptocurrency returns are heavy-tailed and non-Gaussian. Bitcoin's return distribution shows stronger non-normality and heavier tails than traditional currencies (Osterrieder and Lorenz 2017). Extreme value estimates of Value-at-Risk and Expected Shortfall show that tail-risk rankings differ across major cryptocurrencies (Gkillas and Katsiampa 2018). Cryptocurrencies are also exposed to conditional tail risk within crypto markets, even if links to traditional global assets are weaker (Borri 2019).

The implication is that standard deviation is an incomplete risk statistic. It treats upside and downside deviations symmetrically, while investors are typically more concerned with crash losses, liquidity stress, and contagion. Value-at-Risk (VaR) improves the focus on downside quantiles, but it can still understate risk when the distribution is very heavy-tailed or when losses beyond the quantile are severe. Expected Shortfall is more informative because it averages losses conditional on being in the tail, but it also depends on the quality of the tail model and the availability of enough extreme observations. In short cryptocurrency samples, the relevant tail may be defined by only a few crisis events.

The Terra/LUNA collapse is an important example. Algorithmic stablecoins attempt to maintain a peg through arbitrage incentives and endogenous token supply rather than through fully reserved collateral. This design is fragile because it depends on continuing demand, credible incentives, and market confidence (Clements 2021). The UST crash has been modeled as a gradual run with suspension risk (Uhlig 2022), while token-economic incentive failures have also been identified as central to the Terra mechanism (Cho 2023). The common lesson is that a design marketed as stable can embed nonlinear downside risk: once confidence breaks, redemption pressure, supply expansion, and falling token prices reinforce each other.

For portfolio analysis, Terra/LUNA shows why tail events are not merely outliers to be smoothed away. They can change the economic identity of an asset, destroy liquidity, and contaminate estimated moments for years after the event. A strategy that appears diversified before such a collapse may become exposed to a single protocol failure. This is why the thesis treats higher moments, drawdowns, and the LUNA case study as essential complements to mean–variance performance metrics.

2.5 Transaction costs and implementation in portfolio theory

Frictionless portfolio theory assumes that target weights can be implemented immediately and costlessly. In practice, rebalancing costs wealth through commissions, bid–ask spreads, price impact, slippage, and execution delay. These frictions are not secondary details, because the same optimization rule that looks attractive before costs may require frequent trading or large position changes. The economic object of interest is therefore not only the target portfolio $\mathbf{w}_t^*$, but also the trade vector required to reach it:

$$\Delta \mathbf{w}_t = \mathbf{w}_t^* - \mathbf{w}_t^-, \tag{2.4}$$

where $\mathbf{w}_t^-$ denotes the portfolio weights immediately before rebalancing. A proportional-cost approximation writes implementation cost as

$$C_t = \sum_{i=1}^{N} c_{i,t} |\Delta w_{i,t}|, \tag{2.5}$$

where $c_{i,t}$ is the cost rate for asset i. This expression links turnover directly to performance erosion.

The theoretical literature shows that transaction costs change the optimal policy itself. With proportional transaction costs, investors should generally tolerate a no-trade region around the frictionless optimum rather than rebalance continuously (Davis and Norman 1990). This result is conceptually important for empirical backtests: if a strategy trades aggressively to restore exact weights, the reported frictionless performance is not the implementable optimum once costs are present. More recent dynamic models similarly emphasize that optimal portfolios must trade off diversification gains against the cost of moving capital (Gârleanu and Pedersen 2016).

Execution-cost studies also distinguish quoted prices from realized trading outcomes. Institutional trades affect prices around execution (Chan and Lakonishok 1995), while optimal-execution models formalize the trade-off between execution speed, price impact, and risk (Almgren and Chriss 2001). Implementation shortfall measures the difference between the theoretical decision price and the realized execution price, combining explicit fees with implicit costs such as slippage and market impact. This concept is especially relevant for cryptocurrency markets, where liquidity can be fragmented across venues and can deteriorate during stress.

For empirical portfolio research, the practical consequence is clear: turnover is a risk channel. Strategies such as maximum Sharpe or maximum diversification can generate attractive in-sample allocations but may require large rebalances when inputs change. Risk-parity and minimum-variance methods may be more stable, but they are not costless either. A realistic assessment must therefore compare pre-cost and post-cost performance, identify which strategies are most turnover-intensive, and test whether diversification benefits survive implementation. This provides the theoretical basis for the transaction-cost backtest in Chapter 6.

The following chapters operationalize these theoretical foundations as follows: Chapter 3 constructs the cryptocurrency and traditional-asset panels needed to estimate returns, risks, and universe membership without look-ahead bias; Chapter 4 translates mean–variance and Smart Beta 2.0 theory into explicit allocation rules and performance metrics; Chapter 5 evaluates the resulting strategies under heavy-tailed returns and the Terra/LUNA episode; and Chapter 6 tests whether the same allocations remain economically meaningful after transaction costs, slippage, and turnover are introduced.

3 Data collection, selection and description

3.1 Introduction

Data collection is easy to underestimate and surprisingly hard to get right. In cryptocurrency research this is especially true: past work in the area has been hurt repeatedly by selection bias and inadequate historical coverage. Addressing those problems was a primary concern here. CoinGecko was chosen as the source for cryptocurrency data because it provides detailed historical records alongside rich metadata, including category labels, blockchain information, and textual descriptions for each asset. The sections that follow explain how the data were collected and processed, what was selected and why, and what the final dataset looks like.

Traditional financial markets offer a very different starting point from the cryptocurrency space: well-established asset classes, strong regulation, and deep liquidity. To ground the cryptocurrency allocation analysis in a familiar reference, a representative set of traditional assets was included alongside the crypto data. These instruments provide the benchmark exposures against which the risk, return, and diversification properties of cryptocurrencies can be assessed.

The goal was not to replicate a complex institutional portfolio. It was to select a small, highly liquid set of investable instruments that proxy for the core asset classes: equities, fixed income, and commodities. The subsections below describe the data source, the instruments chosen, the extraction methodology, and the selection rationale.

3.2 Cryptocurrency data

3.2.1 Source

CoinGecko was selected as the main data source due to its broad coverage of CCs and the availability of both rich metadata and historical market data. However, due to severe API limitations regarding historical depth, data collection was performed through custom web scraping rather than official endpoints.

3.2.2 Access and extraction methodology

The extraction pipeline was implemented in Python using the `botasaurus` library, which emulates real browser behavior to bypass protections like Cloudflare[1].

Two-stage scraping

- **Metadata scraping:** A script parsed the CoinGecko market rankings to gather all relevant asset metadata including name, ticker, blockchain information, CoinGecko categories, and textual descriptions.[2]
- **Time series scraping:** A second script visited each asset page and downloaded its daily historical data in CSV format including price, market capitalization, trading volume.

Deduplication and file resolution

Due to CCs nature, they can co-exist on different blockchains but have the same price, it was common to have multiple files for the same asset (e.g., `coin.csv`, `coin (1).csv`). To standardize:

- files were grouped by base name;
- the version with the most data rows was retained and renamed canonically;

[1]Cloudflare is a security and performance service that sits between a website and its visitors, filtering traffic and blocking bots. It often uses challenges like JavaScript computations or CAPTCHA to prevent automated access.

[2]CoinGecko asset metadata, accessed 23 May 2025 *CoinGecko* 2025.

- a log file (`duplicates_stats.txt`) recorded all duplicates and selections.

Mapping metadata to filenames

The initial metadata scraping from CoinGecko failed to retrieve the actual filenames of the time series CSVs due to the platform's non-transparent file naming. As a result, a secondary step was required to ensure correct linkage between each CC asset and its corresponding data file. This was handled by a custom script (`mapping.py`), which systematically queried CoinGecko's export URLs and extracted the actual filenames. The script used a robust regular expression parser with fallback mechanisms and sanitized the extracted filenames to ensure safe and consistent file references. The final output was a JSON dictionary (`mapping.json`) establishing a one-to-one mapping between each asset's relative link and its definitive filename, crucial for unambiguous analysis. This simplified access to historical data via the metadata CSV file which is shown later.

Categorization

CoinGecko's native category labels are often multiple and not reliable. A Large Language Model (Gemini 1.5 Flash Latest[3]) was therefore used to classify each CC into one of the following macro-categories:

- Infrastructure & Interoperability - Base-layer protocols and cross-chain systems enabling other applications to function and connect.
- DeFi Primitives - Core decentralized finance building blocks such as lending, borrowing, and exchanges.
- Stablecoins - Tokens pegged to fiat or other assets for price stability in transactions and contracts.
- Liquid Staking & Restaking - Mechanisms for staking assets while retaining liquidity through derivative tokens.
- Bridged & Wrapped Assets - Representations of tokens moved across chains or wrapped into alternative formats.
- NFTs, Gaming & Metaverse - Digital collectibles, game economies, and virtual worlds built on blockchain.

[3]The classification script used the `gemini-1.5-flash-latest` model endpoint. Gemini 1.5 Flash was available at the time of the May 2025 data work but has since been discontinued by Google and replaced by the Gemini 2.x family, which matters for exact reproducibility of the LLM-based classification step.

- AI, Data & IoT - Integration of machine learning, decentralized data markets, and connected devices with blockchain.
- Social, Media & Community - Platforms for decentralized communication, content, and community governance.
- Meme & Culture - Tokens driven by internet culture, social movements, or community trends.
- Real-World Assets - On-chain representations of physical or traditional financial assets.
- Privacy & Security - Protocols focused on confidential transactions, secure computation, and user protection.
- CeFi, Launchpads & VC - Centralized finance services, fundraising platforms, and venture investment structures.
- Other / Unclassified - Assets that do not fit into established categories or span multiple sectors.

Where LLM inference failed (e.g., no usable description), a fallback keyword-based classifier was applied using regex matches on tags and names.

3.2.3 Data specifics

The final dataset consists of:

- A cleaned and enriched metadata file: `coingecko_processed_llm.csv`, produced after deduplication, filename mapping, and LLM-based categorization.
- One individual CSV file per coin, storing its historical daily market data.

The file `coingecko_processed_llm.csv` is not a raw export from CoinGecko. It is the result of a structured processing pipeline that:

1. scraped raw metadata from CoinGecko's market ranking pages;
2. mapped each asset to its corresponding historical CSV;
3. removed duplicates and inconsistencies;
4. applied LLM-driven macro-category classification.

This enrichment was necessary to standardize the dataset and enable sector-level filtering and analysis. In the following the structure of the them is shown.

Metadata dataset structure (`coingecko_processed_llm.csv`)

Column	Description
`Name`	Full name of the cryptocurrency
`Ticker`	Abbreviated symbol
`Filename`	Pointer to historical data CSV file
`Blockchains_JSON`	List of blockchains (JSON format)
`Macro_Category_LLM`	Category assigned by LLM
`Categories_JSON`	CoinGecko-provided tags (JSON format)
`Description_JSON`	Long-form description (JSON format)

Table 3.1: Metadata file structure

Historical data structure

Column	Description
`snapped_at`	Timestamp of daily snapshot
`price`	Closing price in USD
`market_cap`	Market capitalization in USD
`total_volume`	Total daily traded volume in USD

Table 3.2: Daily time series file structure

3.2.4 Limitations

Although the scraped dataset nominally begins on January 3, 2009 (Bitcoin's genesis block [4]), the earliest reliable data from CoinGecko starts on April 28, 2013. This reflects when CoinGecko began structured archival of market prices, not actual coin launch or exchange launch dates. Furthermore, not all assets listed on CoinGecko had historical data available for download. This discrepancy stems from the platform's incomplete archival, not from scraping failure.

The Table 3.3 summarizes data loss across three stages of the pipeline. The first row shows the total number of cryptocurrency entries scraped from the website. In the second stage, historical CSV files were downloaded, with 1.04% of entries missing due to failed downloads or missing historical data. The final

[4] The genesis block is the first block of the Bitcoin blockchain, marking the beginning of the network and containing a message criticizing the legacy financial system.

stage involves mapping CSVs to metadata using a JSON file, where 12.21% of entries could not be linked due to filename mismatches, network issues or other exceptions.

Source	Count	Missing	Missing %
Website (scraped)	17,171	—	—
CSV (historical data)	16,993	178	1.04%
JSON (filename mapping)	15,075	2,096	12.21%

Table 3.3: Comparison of CCs listings vs. Usable and available data (as of 23/05/2025)

Lastly, the one coin per file structure, while efficient for modular scraping, imposes an access cost: time series must be manually matched using the metadata's `Filename` column.

3.3 Traditional data

3.3.1 Source

To represent broad, investable exposures in traditional asset classes, a universe of seven highly liquid exchange-traded funds (ETFs) was selected. These funds cover US equities, US government and corporate bonds, key commodity sectors, and gold, all denominated in US dollars.

The primary data source for these instruments was their most liquid listing venue as available through the Refinitiv[5] data feed. For most ETFs, the RIC[6] ends with the `.P` suffix, which denotes a listing on the NYSE Arca exchange in the United States. For the two bond ETFs (GOVT and USIG), we instead used the `.DF` suffix, which in Refinitiv refers to the NASD/FINRA Alternative Display Facility. The `.P` and `.DF` listings were selected for each asset because, among the available alternatives, they contained the most accurate and complete historical data.

The final universe is reported in Table 3.4 using the same descriptive format later adopted for the cryptocurrency master universe.

[5]Refinitiv is a financial market data provider offering real-time and historical pricing, reference data, and analytics for global securities.

[6]A Reuters Instrument Code (RIC) is a ticker-like identifier used in Refinitiv systems to uniquely specify financial instruments and their trading venues.

3.3.2 Access and extraction methodology

Data was retrieved programmatically via the Refinitiv Python library (`eikon`). This raw daily data, which includes non-trading days, forms the basis for any subsequent calendar alignment or resampling required for the analysis.

For each ETF, the following core Refinitiv fields were extracted:

- `TR.PriceClose` – The daily official closing price.
- `TR.FundTotalNetAssets` – The fund's total net assets (AUM).
- `TR.ACCUMULATEDVOLUME` – The consolidated daily trading volume.
- `TR.SharesOutstanding` – The total number of shares issued by the fund.

The final, processed data is written to per-RIC CSV outputs for auditability and reproducibility.

3.3.3 Data specifics

Table 3.4 summarizes the identifiers and roles of the selected ETFs. Prices are denominated in USD, volumes are consolidated daily totals, and AUM is as reported by Refinitiv on each date.

Symbol	Name	Class / category	Role in analysis
`SPY.P`	SPDR S&P 500 ETF Trust	US equities	Broad US large-cap equity benchmark.
`GOVT.DF`	iShares U.S. Treasury Bond ETF	US treasuries	Government-bond exposure.
`USIG.DF`	iShares Broad USD Investment Grade Corporate Bond ETF	US investment-grade corporates	Corporate-bond exposure.
`DBE.P`	Invesco DB Energy Fund	Commodities	Energy-futures proxy.
`DBB.P`	Invesco DB Base Metals Fund	Commodities	Industrial-metals futures proxy.
`DBA.P`	Invesco DB Agriculture Fund	Commodities	Agricultural-futures proxy.
`GLD.P`	SPDR Gold Shares	Gold	Physically backed gold exposure.

Table 3.4: Traditional asset universe in the common descriptive format.

The following data dictionary describes the final columns available in the output files after processing.

Column	Description
price	Daily close price in USD.
market_cap	Proxy for market capitalization (price × shares_outstanding).
total_volume	Consolidated daily trading volume in shares, forward-filled to create a continuous series.
aum	Total net assets (AUM), forward-filled to create a continuous series.
shares_outstanding	Number of shares, forward-filled to create a continuous series.

Table 3.5: Final dataset columns derived from Refinitiv fields.

3.3.4 Limitations

- **Sparse data and forward filling.** As already hinted some fields are not reported every day. To build continuous daily series the last available value is carried forward, this is standard practice but still an imputation.
- **ETF tracking and structure.** ETF prices do not perfectly track their reference indices because of fees, replication methods, and other frictions. Futures-based commodity funds (DBE, DBB, DBA) also reflect roll yield and collateral effects, so they are proxies for commodities rather than spot prices.
- **Trading calendar alignment.** The raw data includes all calendar days. Weekends and holidays appear with missing price and volume. Any analysis must align ETF trading days with the 24/7 cryptocurrency calendar.
- **Data availability and quality.** All selected ETFs existed before 2015, but market closures or vendor outages can still cause missing values. Gaps in the CSV files reflect what was available from the data provider for those dates.

3.4 Data processing pipeline

The two raw datasets described in Sections 3.3 and the preceding cryptocurrency section are not fed directly into the asset-allocation backtests of Chapters 4 and 5. They first traverse a deterministic, multi-stage pipeline whose purpose is to (i) consolidate the raw vendor outputs into a common schema, (ii) enforce strict temporal discipline before any look-back-dependent statistic is computed, and (iii) produce the per-asset return panels that the optimisers consume. This section narrates the pipeline end-to-end at the methodological level.

Figure 3.1 summarizes the full acquisition, cleaning, classification, alignment, and return-construction workflow.

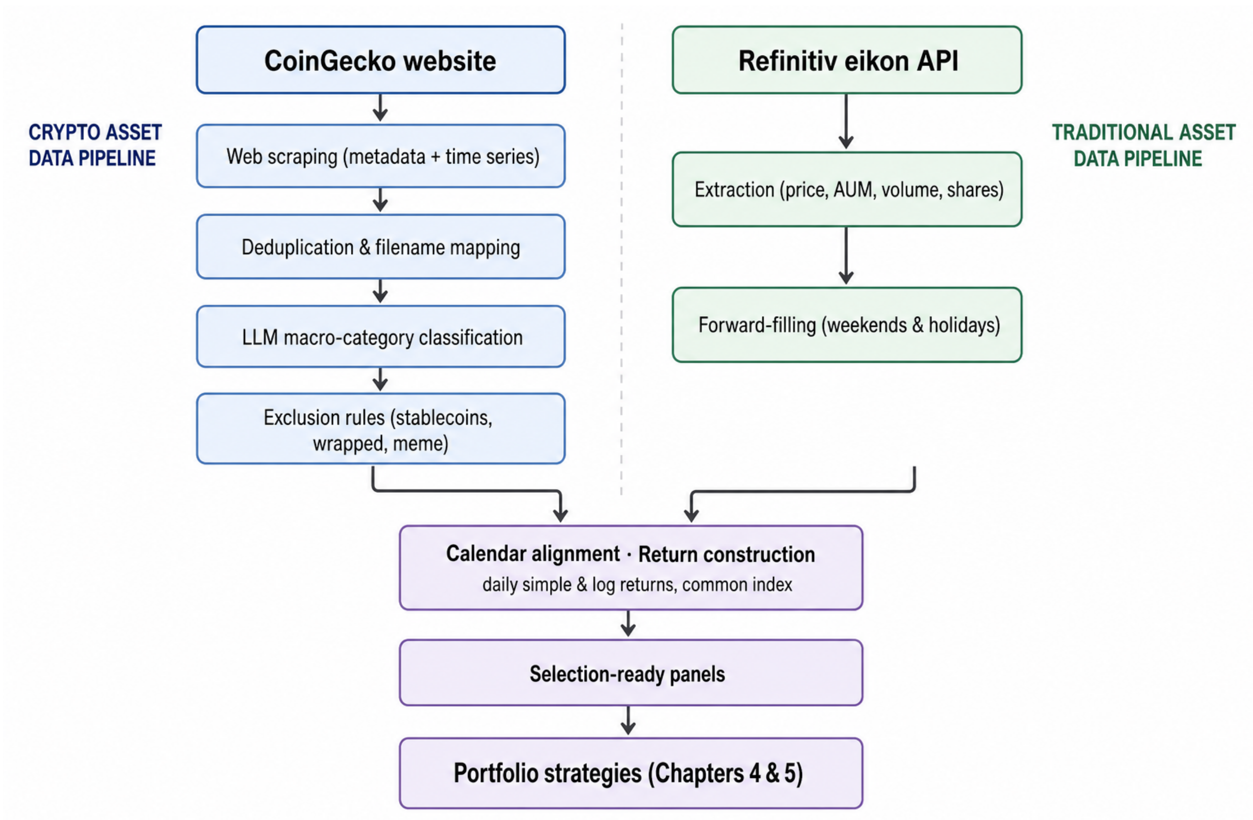

Figure 3.1: End-to-end data acquisition and processing pipeline for the cryptocurrency and traditional asset datasets. Left branch: CoinGecko web scraping, deduplication, filename mapping, and LLM-based macro-category classification. Right branch: Refinitiv `eikon` API extraction and forward-filling. Both branches merge into the joint alignment and return-construction stage that feeds Chapters 4 and 5.

3.4.1 Acquisition layer

The acquisition layer is split by data source. For cryptocurrencies, it is implemented as a two-stage browser-emulating scraper of CoinGecko (Section 3.2.2): a metadata pass collecting names, tickers, blockchains and textual descriptions, followed by a per-asset historical pass that downloads daily `snapped_at`, `price`, `market_cap` and `total_volume` fields into one CSV per coin. For the traditional sleeve, the acquisition layer is a Python client that calls the Refinitiv `eikon` API (Section 3.3) and extracts `TR.PriceClose`, `TR.FundTotalNetAssets`, `TR.ACCUMULATEDVOLUME` and `TR.SharesOutstanding` per RIC into a per-instrument CSV that mirrors the crypto schema as closely as possible. Both branches write timestamped, idempotent outputs to disk so that downstream stages can be rerun without ever hitting the vendor APIs again. As a deliberate methodological choice these acquisition scripts are treated as *frozen*: they are not re-executed once a snapshot has been signed off, to guarantee that the empirical results are reproducible from the on-disk artefacts alone. CoinGecko data was last refreshed on 23 May 2025 and the Refinitiv extraction on the same week.

3.4.2 Cleaning and consolidation

Once acquired, the per-asset CSVs are passed through a cleaning stage that performs four operations in a fixed order. First, duplicate filename resolution: when CoinGecko returns multiple files for the same logical asset (e.g. `coin.csv` and `coin (1).csv`, often due to multi-chain listings), they are grouped by base name and only the version with the largest non-null history is retained, with all decisions journalled to `duplicates_stats.txt`. Second, filename–metadata mapping: the `mapping.py` script queries CoinGecko's export endpoints, parses the returned filenames with a robust regular-expression-based parser and writes a one-to-one correspondence into `mapping.json`, so that every row of the metadata table can be linked unambiguously to a historical CSV. Third, LLM-driven categorisation: each asset's textual description is sent to Gemini 1.5 Flash Latest[7] (with a regex-based fallback for missing or empty descriptions) and assigned to one of the thirteen macro-categories listed in Section 3.2.2; the chosen label is written to the `Macro_Category_LLM` column of `coingecko_processed_llm.csv`. Fourth, exclusion rules: stablecoins, wrapped/bridged assets, liquid-staking tokens and meme tokens are dropped from the eligible pool, because their economic mechanics violate the long-only price-discovery assumptions of the allocation strategies. The analogous traditional-side stage is much simpler: forward-filling of non-trading-day fields and conversion of share-level quan-

[7]See footnote 3 in Section 3.2.2 for the model version and a short reproducibility note.

tities into USD-denominated levels, performed inside the Refinitiv extraction script.

3.4.3 Alignment and return construction

The cleaned per-asset price series are then aligned onto a common daily index. Because cryptocurrency markets trade 24/7 while ETFs trade only on US business days, the alignment policy is asymmetric. ETF closing prices are forward-filled across crypto-only trading days (weekends, US holidays) to preserve a continuous reference path; this is standard practice in mixed crypto–equity studies and matches the convention used for the risk-free series in Chapter 5 (Campbell, Lo, and MacKinlay 1997). Cryptocurrency prices, by contrast, are taken at face value on every calendar day. Daily simple returns are then computed as $r_t = p_t/p_{t-1} - 1$ for each asset, and log-returns as $\ell_t = \log(p_t/p_{t-1})$ are stored alongside for the distributional diagnostics of Section 3.6. No outright outlier removal is performed on the return series: empirical evidence in Section 3.6 (and in the cryptocurrency literature more broadly) shows that extreme observations are a genuine feature of the data-generating process, not a data-quality artefact, and removing them would understate tail risk. The boxplots of Figure 3.6 and Figure 3.8 nevertheless display both the full and the $1.5\times$IQR-trimmed distributions side-by-side, so that the contribution of extremes to each asset's profile remains visible.

3.4.4 Selection-ready panels and reproducibility

The output of the pipeline is a set of selection-ready panels: a metadata table with LLM-assigned categories, a long-format date $\times$ asset price and market-cap matrix for the eligible crypto universe, and a fixed seven-instrument matrix for the traditional sleeve. These panels are the single point of entry for the rolling selection rules of Section 3.5 and for every allocation strategy implemented in Chapter 4. The acquisition, cleaning, and alignment layer is version-controlled and treated as frozen: figures in this chapter are regenerated only from the on-disk panels, never by re-querying vendor APIs. This separation makes the empirical results in Chapters 5 and 6 reproducible from the committed artefacts.

3.5 Selection criteria for asset allocation

3.5.1 Cryptocurrency data

To ensure economic relevance and avoid distortions caused by micro-cap, illiquid, or defunct tokens, a filtering process was applied based on historical relative market capitalization. The filtering covers the period from **1st January 2015 to 31st December 2023**, with monthly rebalancing.

The procedure was designed with strict temporal discipline and investment realism in mind. First, it is applied at the macro-category level to identify which macro-category contains the largest number of cryptocurrencies meeting the criteria, thereby removing noise from the full set of 15,075 cryptocurrencies. The same procedure is then applied within the selected macro-category to obtain the individual cryptocurrencies used to build the investment universe. Key filtering rules:

1. **Rolling selection window:** At each portfolio rebalancing date T, only data from the backward-looking window $[T-M,\ T-1]$ is used. No information from time T or beyond is accessed, to strictly prevent look-ahead bias.

2. **Coin eligibility:** A coin must be *continuously active* throughout the entire window. Any coin that is created or delisted within the period $[T-M,\ T-1]$ is excluded from selection.

3. **Market cap snapshot:** Filtering is based solely on the market capitalization at the final day of the window, i.e., at $T-1$. No averaging or smoothing is performed.

4. **Relative share calculation:** For each macro-category C, the relative market share at time $T-1$ is calculated as:

$$\text{Relative Share}_C(T-1) = \frac{\sum\limits_{i \in C} \text{Market Cap}_i(T-1)}{\sum\limits_{j} \text{Market Cap}_j(T-1)}$$

   where the numerator sums the market capitalizations of all coins i in category C, and the denominator includes all eligible coins on that day.

5. **Thresholding rule:** Categories with a relative share below a fixed threshold (e.g., 1% or 5%) are excluded from the investable universe at that rebalancing date.

6. **Rebalancing:** The filtering is repeated at each month-end. The window advances forward one month at a time to dynamically reflect the evolving market structure.

This methodology ensures that only economically significant and liquid macro-categories are retained, while enforcing a clean out-of-sample setup free from survivorship and forward-looking bias.

3.5.1.1 Rationale for the parameter choices

The two free parameters of the selection rule, namely the relative market-capitalisation threshold and the length of the look-back window, are not picked arbitrarily: each addresses a specific empirical concern raised during the supervision of this thesis and each can be defended on independent economic and statistical grounds. The treatment of survivorship is the third, implicit, methodological choice that completes the rule and is therefore discussed alongside the other two.

1% relative market-capitalisation threshold. The threshold is defined within the selected macro-category so that it remains meaningful as the category's total market capitalization changes over time. At 1%, the eligible set remains liquid and economically material in the sense discussed by Liu and Tsyvinski (2021b), while preserving enough names for portfolio estimation; Figure 3.3 and Figure 3.4 show that a 5% threshold would collapse the universe into a much narrower set.

365-day rolling window. A one-year window covers a full market cycle while remaining responsive to structural breaks in the cryptocurrency market. The window is rolling rather than expanding, so each monthly rebalance uses only the latest information and remains comparable to the design in Petukhina et al. (2021).

Survivorship bias. The full historical CoinGecko download is retained rather than only currently listed coins, and the filter is applied *ex ante* at each rebalance. A coin is eligible only if it traded continuously during the preceding 365 days; the only discretionary intervention is the post-May 2022 exclusion of Terra Classic (LUNC), whose treatment is discussed in Appendix B.

Figure 3.2: Relative market capitalization of CCs macro-categories

For the empirical analysis that follows, only the macro-category **Infrastructure & Interoperability** is retained for selection. As evident from Figure 3.2, this category consistently dominates the market, accounting for over 90% of total market capitalization throughout most of the observed period. In recent years, its dominance slightly decreased to approximately 70%, primarily due to the growing share of stablecoins. However, stablecoins are excluded from this analysis due to their structural characteristics (price anchoring, limited volatility, and non-investable dynamics), which make them unsuitable for the type of allocation strategies considered in this study.

Applying the rules above (specifically the survivorship and 1% market cap threshold within the selected category) at each month-end yields a **dynamic investable universe**. The list below represents the **master universe** of all 21 cryptocurrencies that were eligible for portfolio inclusion at least once during the study period.

Table 3.6: Cryptocurrency master universe in the common descriptive format.

Symbol	Name
ADA	Cardano
AVAX	Avalanche
BCH	Bitcoin Cash
BNB	BNB
BTC	Bitcoin
BSV	Bitcoin SV
CRO	Cronos
DOT	Polkadot
EOS	EOS
ETC	Ethereum Classic
ETH	Ethereum
IOTA	IOTA
LINK	Chainlink
LTC	Litecoin
LUNC	Terra Luna Classic
MATIC	Polygon
NEO	NEO
SOL	Solana
TRX	TRON
XLM	Stellar
XRP	XRP

As already mentioned, it is crucial to understand that this is not a static portfolio where all assets are included simultaneously. Rather, it is the superset from which a smaller, time-varying basket of assets is selected each month.

This dynamic selection process is visualized in the Gantt style plots of Figure 3.3 and Figure 3.4. The horizontal axis represents the rebalancing dates (month-ends), and the vertical axis lists the tickers from the master universe. A black cell signifies that a token was part of the eligible set at that specific month-end,

having met all inclusion criteria.
While the 5% threshold (Figure 3.4) isolates only the most dominant assets, the 1% threshold (Figure 3.3) provides a broader, more inclusive view of the evolving ecosystem. For this reason, the 1% threshold is adopted for the remainder of the analysis, as it yields a richer cross-section of relevant tokens while still enforcing the strict survivorship and liquidity filters.

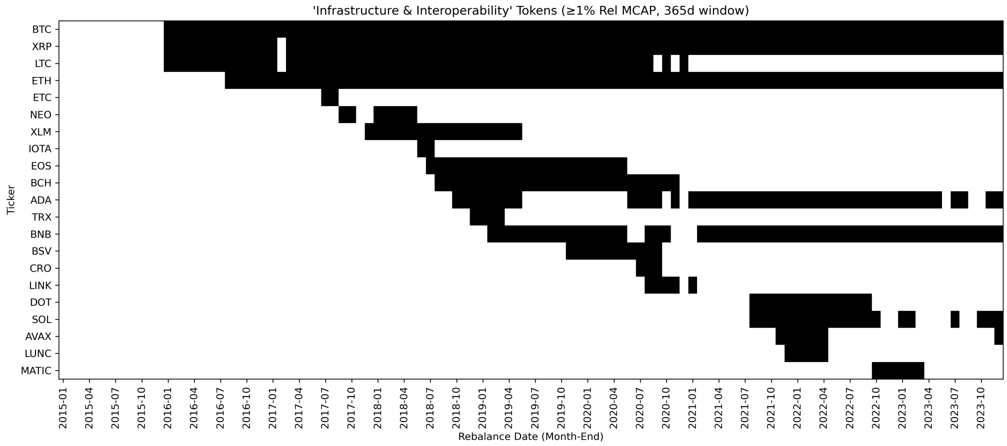

Figure 3.3: 'Infrastructure & Interoperability' Tokens (≥1% Relative market capitalization, 365-day window)

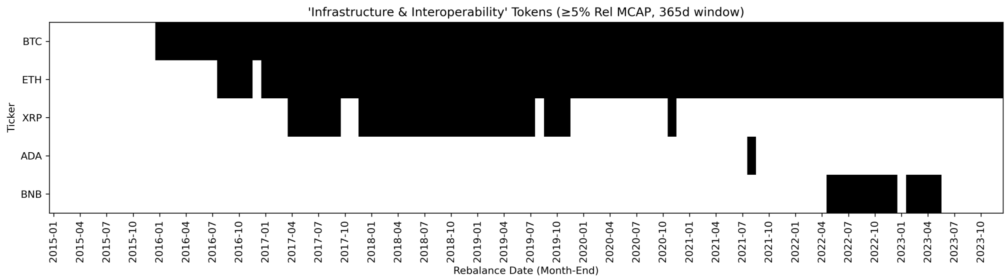

Figure 3.4: 'Infrastructure & Interoperability' Tokens (≥5% Relative market capitalization, 365-day window)

3.5.2 Traditional data

The selection of traditional assets was guided by a principle of stable representation rather than dynamic filtering. Unlike the cryptocurrency universe, which requires monthly re-evaluation to capture its rapidly evolving composition, the traditional asset classes are mature and well-defined. The objective was to establish a fixed, high-quality benchmark universe that remains constant throughout the 2015–2023 analysis period.

The core selection criteria were:

1. **Asset class representation:** The chosen instruments needed to provide broad, diversified exposure to the most fundamental components of a multi-asset portfolio: US large-cap equities, US government bonds, US investment-grade corporate bonds, a diversified basket of commodities, and gold.

2. **Investability and liquidity:** Exchange-Traded Funds (ETFs) were selected as the exclusive investment vehicle. ETFs offer intra-day liquidity, low transaction costs, and are directly accessible to a wide range of investors. The specific ETFs chosen are leaders in their respective categories, commanding billions in assets under management (AUM) and high daily trading volumes, which ensures that their pricing is efficient and reliable.

3. **Static universe:** The set of seven ETFs is held constant for the entire study period. This is a deliberate methodological choice since they are a de-facto proxies for their respective asset classes. A dynamic selection process was deemed unnecessary and counterproductive, as it would introduce needless complexity without altering the fundamental exposures.

This curated, static set of seven ETFs thus forms the **traditional investable universe**. It provides a stable and consistent frame of reference against which the performance and characteristics of the dynamic cryptocurrency portfolio can be rigorously assessed.

3.6 Descriptive analysis

3.6.1 Cryptocurrency data

Table 3.7 reports the key descriptive statistics for daily arithmetic returns of the selected cryptocurrencies. To make the cryptocurrency and traditional-asset summaries directly comparable, returns are reported in percentage points in both tables. The metrics cover central tendency, dispersion, distributional asymmetry, and tail thickness. The results highlight the strongly non-normal nature of cryptocurrency returns, with several assets displaying large skewness and excess kurtosis. For a detailed inspection of residual patterns, volatility clustering, and serial dependence, see Appendix A.

Asset	Min	Q1 (25%)	Median	Mean	Q3 (75%)	Max	Std Dev	Skewness	Kurtosis
ADA	-40.8100	-2.6400	0.0400	0.3400	2.5900	139.2100	6.8600	5.2429	90.1144
AVAX	-35.0100	-3.0300	0.0700	0.4000	3.2800	75.7000	7.0100	1.5613	14.7843
BCH	-48.0500	-2.5200	-0.0400	0.1800	2.4300	139.8200	7.0700	4.0183	73.7037
BNB	-63.3000	-1.9100	0.1100	1.6000	2.4800	2530.9700	53.4000	46.3433	2195.3164
BSV	-47.4300	-2.2900	-0.0500	0.2700	2.0400	142.6900	7.6300	5.3935	84.4226
BTC	-35.1900	-1.2100	0.1600	0.2200	1.7200	33.2600	3.7200	-0.0566	8.2672
CRO	-39.9700	-2.1400	0.0500	0.2700	2.1700	122.8400	6.4200	5.1670	87.8085
DOT	-37.3600	-2.7500	-0.0300	0.2600	2.6500	56.0400	5.9600	1.2445	11.6716
EOS	-38.6600	-2.3500	-0.0200	0.1600	2.4400	52.5800	6.4400	0.9513	9.1990
ETC	-41.3600	-2.4400	-0.0300	0.3700	2.4100	266.2500	8.1000	13.4651	429.3754
ETH	-53.0000	-2.0600	0.0600	0.3900	2.6900	55.2400	5.8000	0.2598	11.7086
IOTA	-42.9500	-2.8000	0.0500	0.1900	2.8500	50.4300	6.5100	0.8082	9.0795
LINK	-48.3600	-3.1800	0.0900	0.4100	3.3500	60.9700	6.8000	0.9797	9.3107
LTC	-42.1400	-2.0500	0.0300	0.2500	2.1700	67.2600	5.5500	1.2638	16.3679
LUNC	-99.8900	-3.0100	-0.3100	0.3900	2.6700	286.6600	11.8400	8.1621	212.6358
MATIC	-50.3800	-3.1900	0.0100	0.6000	3.2400	68.1000	7.9100	1.7400	13.7405
NEO	-57.6600	-2.8800	0.0600	0.4500	2.9800	205.6400	8.9800	7.2551	137.5004
SOL	-42.2000	-3.2100	0.0000	0.6100	3.8300	46.8900	7.2900	0.5292	4.8837
TRX	-42.5600	-2.0700	0.1500	0.4000	2.1500	122.3500	7.3400	5.8089	83.5616
XLM	-65.1300	-2.5800	-0.0900	0.4700	2.4200	411.8300	10.8400	18.1272	644.1599
XRP	-42.2800	-1.9900	-0.1200	0.3000	1.8700	141.4000	6.7600	4.9936	77.4856

Table 3.7: Descriptive statistics of daily arithmetic returns, in percentage points, for selected CCs.

Full CCs names corresponding to these tickers are listed in Table 3.6.

Figure 3.5 shows the estimated probability density functions (PDFs) of daily log-returns for a selection of CCs over the period 2015–2023. The black dashed line represents a fitted normal distribution with pooled parameters $\mu = -0.0010$, $\sigma = 0.0767$, for comparison.

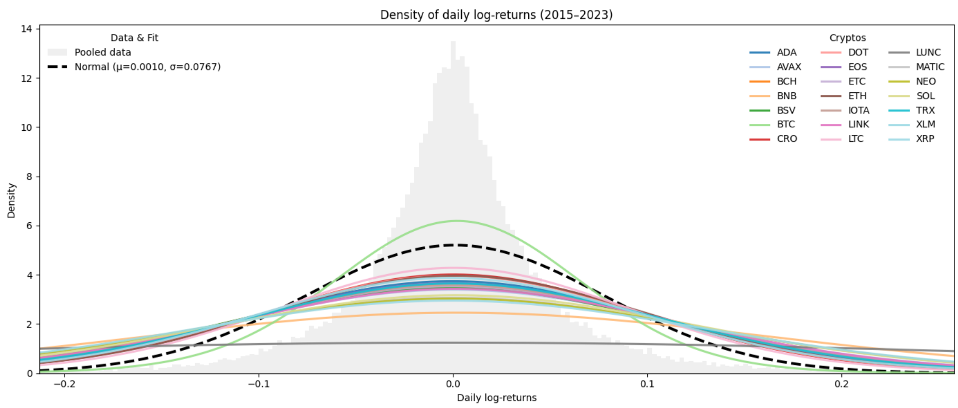

Figure 3.5: Kernel density estimates of daily log-returns for 21 CCs (2015–2023), overlaid with a pooled normal distribution.

The density plots show pronounced deviations from the normal benchmark. Most cryptocurrencies display sharp peaks at zero with extremely heavy tails, producing the strong leptokurtic shapes visible against the Gaussian fit. This is consistent with the high kurtosis estimates reported in Table 3.7. The excess kurtosis reflects a clustering of small daily changes alongside an elevated frequency of large outliers, a pattern characteristic of crypto-asset return dynamics. Across individual tokens, tail thickness varies, but the general pattern of peakedness and fat tails is consistent. These results confirm that Gaussian assumptions are inadequate for modeling crypto returns and underscore the importance of accounting for higher moments when assessing volatility and risk exposure.

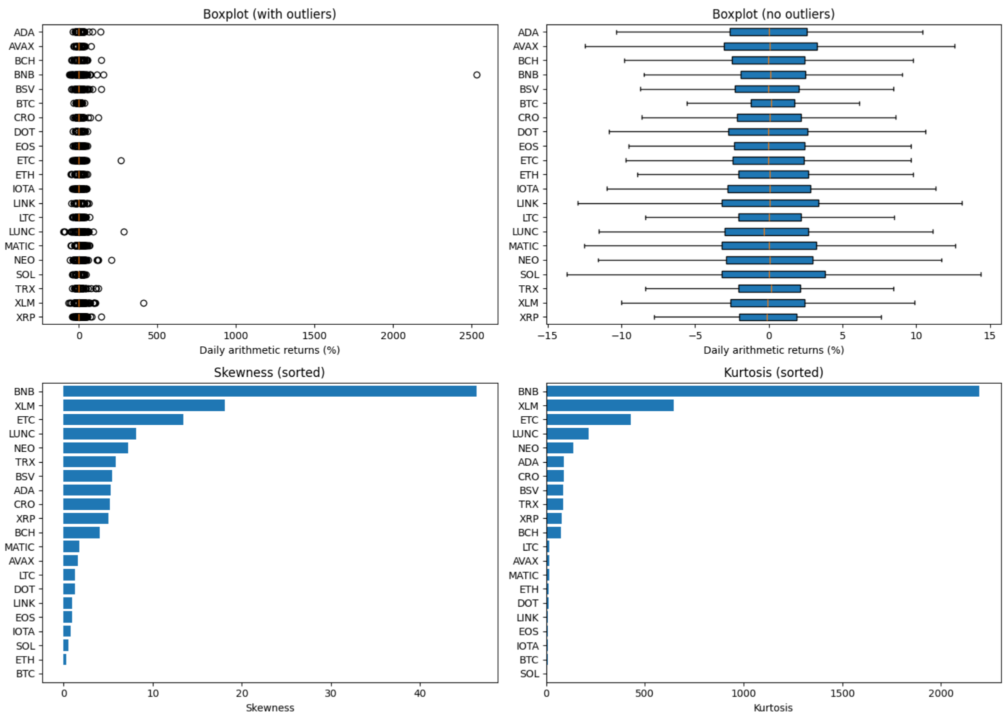

Figure 3.6: Comparison of boxplots for daily arithmetic returns with and without outliers for selected CCs (2015-2023).

Figure 3.6 provides a four-part visual analysis of the daily arithmetic returns for the selected cryptocurrencies, designed to inspect their distributional properties both with and without statistical outliers, skewness, kurtosis.

- The TOP-LEFT PLOT reveals the full extent of return volatility. It shows that several assets have extreme values, most notably BNB, whose isolated jump dominates the scale of the full-sample plot. To a lesser extent, XLM, LUNC, and ETC also exhibit significant, infrequent events. The presence of these isolated points far from the central range underscores the heavy-tailed nature of these return series.

- The TOP-RIGHT PLOT, which excludes extreme values based on the 1.5×IQR criterion, allows for a clearer comparison of core return behavior. Once the outliers are removed, most assets exhibit more symmetric distributions centered around zero, with visually similar interquartile ranges. This view suggests that the bulk of daily price changes are relatively modest and statistically comparable across assets.
- The SKEWNESS PLOT (BOTTOM-LEFT) confirms the asymmetry, ranking assets like BNB, XLM, and ETC highest. Similarly, the KURTOSIS PLOT (BOTTOM-RIGHT) provides a stark visual of the extreme tail risk, with BNB's kurtosis value far exceeding all other assets.

Together, these plots show how a few high-impact observations can distort the representation of typical behavior. Removing them allows for better comparative insights, though it also risks obscuring the tail risk that characterizes CC markets. Both views are therefore necessary: one captures the volatility extremes, while the other represents the normal state.

3.6.2 Traditional data

Table 3.8 reports the key descriptive statistics for daily arithmetic returns of all seven traditional-asset ETFs in the universe of Table 3.4. The table follows the same format and units as the cryptocurrency summary: returns are expressed in percentage points, while skewness and kurtosis are unit-free. The results illustrate the diverse distributional profiles across asset classes. Equities (`SPY`) combine moderate dispersion with negative skewness and heavy tails, energy (`DBE`) has the largest dispersion in the traditional sample, fixed-income proxies (`GOVT`, `USIG`) are concentrated around zero but still display episodic tail risk, and gold (`GLD`) sits between equities and bonds in dispersion and tail thickness. For a detailed inspection of residual patterns, volatility clustering, and serial dependence, see Appendix A.

Asset	Min	Q1 (25%)	Median	Mean	Q3 (75%)	Max	Std Dev	Skewness	Kurtosis
SPY	-10.9424	-0.1529	0.0000	0.0301	0.2949	9.0603	0.9506	-0.6108	19.2866
GOVT	-2.1776	-0.0939	0.0000	-0.0024	0.0863	1.9553	0.2635	0.1273	7.1950
USIG	-4.0291	-0.0894	0.0000	-0.0019	0.1110	5.8481	0.3602	0.4248	43.1444
DBE	-14.0465	-0.4205	0.0000	0.0146	0.5982	8.3952	1.5262	-0.6385	7.4438
DBB	-8.3709	-0.3279	0.0000	0.0092	0.3746	6.3093	0.9866	-0.1984	4.7379
DBA	-4.2805	-0.2723	0.0000	-0.0034	0.2310	2.9757	0.6585	-0.2915	3.8633
GLD	-5.3694	-0.2147	0.0000	0.0185	0.2841	4.9559	0.7329	-0.0226	5.7474

Table 3.8: Descriptive statistics of daily arithmetic returns (in percentage points) for the seven traditional-asset ETFs over 2015–2023. Returns are computed on the resampled, forward-filled daily price series produced by the data-processing pipeline of Section 3.4. The median is zero for all ETFs because forward-filling over weekends and US holidays creates exact zero-return observations on non-trading days after alignment to the daily cryptocurrency calendar. Values are reported to four decimal places; kurtosis is in Fisher (excess) form.

Full asset names corresponding to these tickers are reported in Table 3.4.

Figure 3.7 shows the estimated probability density functions (PDFs) of daily log-returns for the traditional assets. The black dashed line represents a fitted normal distribution with pooled parameters, provided as a benchmark.

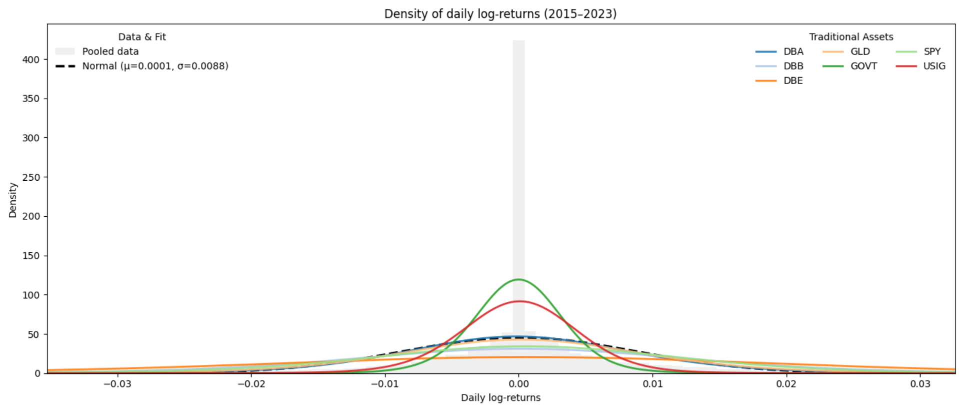

Figure 3.7: Kernel density estimates of daily log-returns for traditional assets, overlaid with a pooled normal distribution.

The empirical densities deviate systematically from normality. Traditional assets have lower dispersion than cryptocurrencies but remain clearly non-Gaussian. Most show excess kurtosis, with sharper peaks, fat tails, and occasional asymmetry. Energy (DBE) displays the heaviest tails, consistent with its high intrinsic volatility, while government bonds (GOVT) are tightly concentrated around zero. Equities (SPY) combine central mass with fat tails, in line with equity risk premia. Commodities (DBA, DBB) and gold (GLD) fall between the two extremes, reflecting intermediate risk–return profiles. Investment-grade credit (USIG) is moderately leptokurtic, balancing stable yield behavior with episodic stress. Overall, the shapes differ across classes but consistently diverge from the Gaussian benchmark.

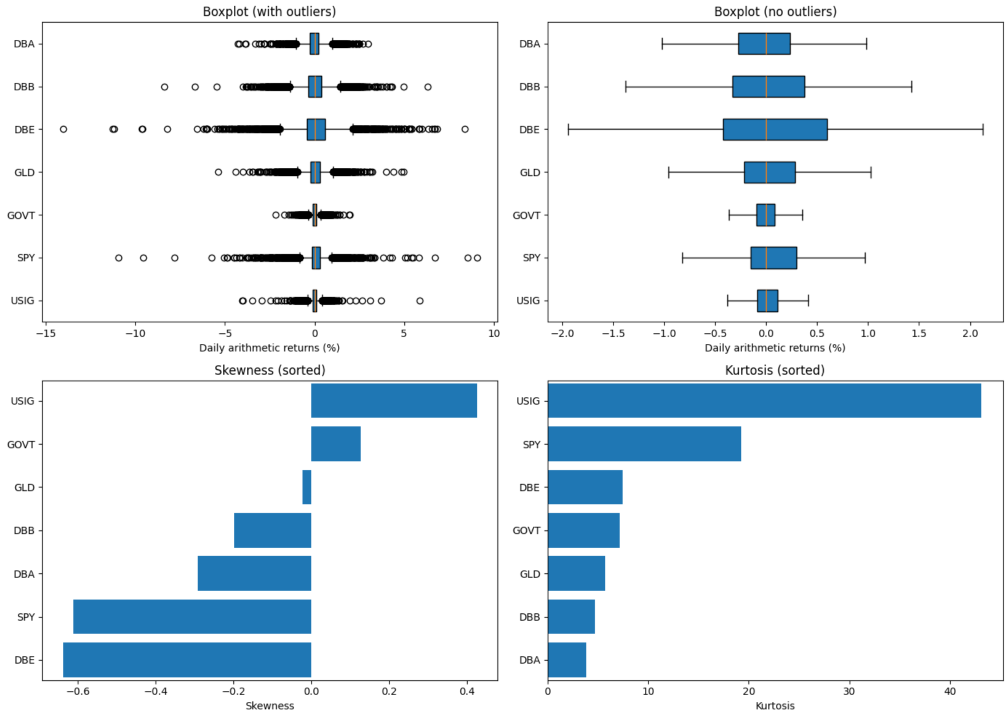

Figure 3.8: Comparison of boxplots for daily arithmetic returns with and without outliers for traditional assets.

Figure 3.8 parallels the cryptocurrency analysis:

- The TOP-LEFT PLOT reveals the presence of extreme values across the traditional-asset universe, with commodity and equity exposures showing the widest excursions. Several observations lie far outside the central range, emphasizing the heavy-tailed behavior of these series.

- The TOP-RIGHT PLOT, adjusted by excluding extreme observations, provides a clearer representation of the core return dynamics. This view highlights the more typical behavior of equities, bonds, commodities, and gold without distortion from outliers.

- The SKEWNESS PLOT (BOTTOM-LEFT) illustrates the degree of asymmetry across assets, while the KURTOSIS PLOT (BOTTOM-RIGHT) highlights the pronounced fat tails of the most stress-sensitive series. In contrast, `GOVT` displays more stable and regular distributional characteristics.

References

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay. 1997. *The Econometrics of Financial Markets*. Princeton, NJ: Princeton University Press.

CoinGecko. 2025. *CoinGecko*. Accessed 23 May 2025. https://www.coingecko.com

Liu, Yukun, and Aleh Tsyvinski. 2021. "Risks and Returns of Cryptocurrency." *The Review of Financial Studies* 34(6): 2689–2727. https://doi.org/10.1093/rfs/hhaa113

Petukhina, Alla, Simon Trimborn, Wolfgang Karl Härdle, and Hermann Elendner. 2021. "Investing with Cryptocurrencies – Evaluating Their Potential for Portfolio Allocation Strategies." *Quantitative Finance* 21(11): 1825–1853. https://doi.org/10.1080/14697688.2021.1880023

4 Asset allocation strategies: methodology and framework

4.1 Introduction

This chapter details the empirical analysis of the portfolio allocation strategies, using a combination of cryptocurrency and traditional assets. Performance is benchmarked against two passive reference portfolios: an equally weighted (EW) portfolio and a value-weighted (VW) portfolio, both constructed from the same set of eligible assets.

The analysis examines two main families of allocation methods. The first is the classical mean-variance framework introduced by Markowitz (1952), which includes three primary strategies. The Global Minimum Variance (GMV) portfolio is constructed based solely on the covariance matrix to find the allocation with the lowest possible variance. The Maximum Sharpe Ratio (MSR) or Tangency portfolio aims for the optimal risk-adjusted return, but its reliance on both mean and covariance estimates makes it sensitive to parameter uncertainty. The framework is generalized by utility-maximizing portfolios, which incorporate an investor's risk-aversion parameter, γ, to tailor allocations to specific risk profiles.

The second group of strategies falls under the Smart Beta 2.0 paradigm, as described by Amenc, Goltz, and Martellini (2013). These risk-based approaches are designed to be more stable than traditional mean-variance optimization. The Maximum Decorrelation (MDC) portfolio minimizes the overall correlation between assets. The Maximum Diversification Ratio (MDR) portfolio seeks to maximize the benefits of diversification by maximizing the ratio of weighted average asset volatility to total portfolio volatility. The Equal Risk Contribution (ERC) portfolio, based on the principle of risk parity, allocates weights so that each asset contributes equally to total portfolio risk. A simpler version, the Diversified Risk Parity (DRP) portfolio, assumes equal correlations and weights assets inversely to their volatilities.

All strategies are implemented with a consistent methodology: parameters are estimated using rolling windows, portfolios are rebalanced monthly, and short-selling is not permitted. Their performance is evaluated using a set of metrics covering return, risk, diversification, and trading costs. Risk-adjusted performance is measured by the Sharpe ratio (Sharpe 1964) and the Adjusted Sharpe ratio, which accounts for higher moments (Pezier and White 2008). Diversification is assessed through the effective number of assets and the portfolio diversification index (Choueifaty and Coignard 2008). Finally, implementation costs and stability are evaluated using turnover and target turnover metrics (Petukhina et al. 2021).

4.2 Methodological framework

This section outlines the universal rules governing the implementation and backtesting of all portfolio strategies. A consistent framework is essential for an unbiased comparison across all models, especially when dealing with the evolving and volatile structure of the cryptocurrency market.

4.2.1 Asset universe construction

The analysis is conducted on two distinct asset universes: a dynamically selected set of cryptocurrencies and a static basket of traditional assets used as a benchmark.

4.2.1.1 Cryptocurrency universe

The cryptocurrency portfolio is composed of assets from the *Infrastructure & Interoperability* macro-category, as defined by a Large Language Model classifier applied to CoinGecko metadata (see Chapter 3). To ensure liquidity, economic relevance, and temporal integrity, the following strict, data-driven selection process was adopted at each monthly rebalancing point:

- **Rolling window**: At each rebalancing date T, only data from the trailing M-month window $[T - M, T - 1]$ is used. This prevents look-ahead bias.
- **Survivorship filter**: Assets must have been continuously active throughout the full look-back window. Coins launched or delisted within this period are excluded from the sample for that specific rebalancing point.
- **Relative market capitalization filter**: Only coins representing at least 1% of the macro-category's total market capitalization at the end of the look-back period ($T - 1$) are deemed eligible for inclusion.
- **Dynamic inclusion**: The resulting list of eligible coins is updated monthly, reflecting the natural evolution of the market.

A single, strategically motivated exception was made to this otherwise purely systematic process. Following its protocol failure and subsequent hyperinflationary collapse in May 2022, Terra Classic (LUNC) was manually excluded from the investable universe for all subsequent periods. This was a necessary intervention to prevent a unique, idiosyncratic asset failure from fundamentally distorting the performance and risk metrics of the portfolio strategies being evaluated.

This filtering pipeline ensures that only investable, liquid, and economically significant assets are included at each step. For the period studied, this procedure yields a dynamic universe of up to 21 cryptocurrencies per month (see Table 3.6 for the complete list).

4.2.1.2 Traditional asset universe

In contrast to the dynamic cryptocurrency universe, the traditional asset portfolio is static. It consists of a fixed set of seven highly liquid, U.S. dollar-denominated exchange-traded funds (ETFs) that remain constant throughout the entire 2015–2023 analysis period. These ETFs were selected to provide broad, investable exposure to core asset classes: U.S. equities, U.S. government and corporate bonds, commodities, and gold. This deliberate choice ensures a stable and consistent frame of reference for benchmarking the cryptocurrency strategies. (see Table 3.4 for the complete list)

4.2.2 Portfolio constraints

To ensure practical feasibility and reflect typical market conditions, all optimised portfolios are subject to two primary constraints:

- **No short-selling**: All portfolio weights must be non-negative ($w_i \geq 0$ for all i).
- **Full investment**: The sum of all weights must equal one ($\sum w_i = 1$).

These constraints are applied to all strategies except for the benchmark portfolios, which are constructed by definition.

4.2.3 Rebalancing protocol and parameter estimation

All portfolios are rebalanced on a monthly basis. At the end of each month $t-1$, the statistical parameters required by the allocation models are estimated using a fixed look-back period of M months via a rolling window. This strict walk-forward procedure simulates a realistic investment process and avoids any contamination from future information.

The key parameters estimated are:

1. **The vector of expected returns (μ)**: The sample mean of historical returns within the window. This is used only by the mean-variance framework strategies.
2. **The covariance matrix of returns (Σ)**: The sample covariance matrix of historical returns. This is used by nearly all optimization strategies, both mean-variance and risk-based. From this, the correlation matrix (Ω) and the vector of volatilities (σ) can be derived.

Once estimated, these parameters are used to compute target portfolio weights, which are implemented and held fixed during the subsequent month t.

4.2.4 Benchmark and base portfolios (EW, VW)

Before evaluating optimised strategies, we establish two foundational benchmarks that serve as the baseline for this analysis. These portfolios are constructed using simple, transparent, and non-predictive rules, meaning they do not rely on the statistical estimation of future market parameters like expected returns or covariances. Their purpose is to provide an unbiased reference point against which the performance of more complex models can be judged. Any sophisticated strategy must first demonstrate its ability to outperform these simple, low-cost alternatives to justify its complexity and potential for estimation error. They represent naive diversification (Equally weighted) and the market consensus (Value weighted), respectively.

4.2.4.1 Equally weighted (EW) portfolio

The equally weighted portfolio assigns an identical weight to each of the N assets in the selected universe. The weight for each asset i is defined as:

$$w_i = \frac{1}{N}$$

Despite its simplicity, the EW strategy has been shown to be a surprisingly difficult benchmark to outperform, particularly when parameter uncertainty is high (DeMiguel, Garlappi, and Uppal 2009), as it completely avoids the estimation errors inherent in optimised models. Its primary weakness, however, is that it ignores all market information, such as an asset's size or volatility, which can lead to an unintentional over-concentration in the riskiest assets.

4.2.4.2 Value-weighted (VW) portfolio

The value-weighted, or market capitalisation weighted (VW), portfolio assigns weights based on an asset's market capitalisation relative to the total. It is considered the 'equilibrium' portfolio under the Capital Asset Pricing Model (CAPM) (Sharpe 1964), representing the aggregate consensus of all market participants. It is worth noting, however, that the foundational assumptions of the CAPM are not fully met in the cryptocurrency space, making the VW portfolio a theoretical benchmark rather than a guaranteed equilibrium state. Its practical drawback in this market is the risk of poor diversification, as high market concentration can lead to a portfolio whose performance is dominated by just one or two constituents.

The weight for each asset i is defined as:

$$w_i = \frac{\mathrm{MC}_i(T-1)}{\sum_{j=1}^{N} \mathrm{MC}_j(T-1)}$$

where $\mathrm{MC}_i(T-1)$ is the market capitalization of asset i at the end of the look-back period $T-1$, and the sum runs over all N eligible assets at that rebalancing date.

4.3 Portfolio optimisation: the Markowitz framework

The foundation of modern portfolio theory, introduced by Markowitz (1952), provides a mathematical framework for assembling portfolios that optimise the trade-off between risk (measured as variance) and expected return. While theoretically elegant, its practical application faces significant challenges. The classical framework assumes investor preferences are fully characterised by the first two moments of asset returns (mean and variance), which is only strictly true for quadratic utility or under the assumption of normally distributed returns. As noted by Michaud (1989), mean-variance optimisers can act as 'error maximisers', amplifying the impact of estimation errors in the inputs and often producing unstable and unintuitive portfolio allocations. These practical limitations motivate the exploration of more stable strategies.

4.3.1 Global minimum variance (GMV) portfolio

The GMV portfolio is the unique portfolio on the efficient frontier that exhibits the lowest possible variance. It is a purely risk-based strategy derived directly from the Markowitz framework. Its primary strength is its complete independence from the vector of expected returns, which is widely considered the most unstable and error-prone input in portfolio optimisation (Chopra and Ziemba 1993). This allows the GMV to focus solely on the diversification benefits offered by the covariance structure. The main drawback, however, is that by ignoring returns entirely, it may select assets with very low volatility but also extremely poor return prospects, potentially sacrificing significant upside. The portfolio is found by solving the following quadratic optimisation problem:

$$\min_{w} w^T \Sigma w \quad \text{s.t.} \quad w^T \mathbf{1} = 1, \; w_i \geq 0$$

4.3.2 Tangency portfolio (Maximum Sharpe Ratio)

The tangency portfolio, also known as the Maximum Sharpe Ratio (MSR) portfolio, represents the theoretical optimum for an investor, as it maximises the expected excess return per unit of risk. This theoretical optimality is also its greatest practical weakness. The MSR portfolio requires estimates for both the expected return vector (μ) and the covariance matrix (Σ), making it highly susceptible to parameter estimation risk. Small errors in the estimation of μ, in particular, can lead to extreme and unstable portfolio weights, often resulting in poor out-of-sample performance compared to simpler benchmarks (DeMiguel, Garlappi, and Uppal 2009). At each monthly rebalancing date t, the optimisation is solved on the in-sample window $[t-M, t-1]$:

$$\max_{w} \frac{w^T \mu - r_{f,t}}{\sqrt{w^T \Sigma w}} \quad \text{s.t.} \quad w^T \mathbf{1} = 1,\ w_i \geq 0$$

where μ is the vector of estimated expected returns over the look-back window and $r_{f,t}$ is the risk-free rate prevailing at $t-1$. Equivalently, when the optimisation is iterated forward across the full backtest, r_f is treated as a vector $r_f = (r_{f,t_1}, r_{f,t_2}, \ldots, r_{f,t_T})^T \in \mathbb{R}^T$ of period-specific values rather than a single scalar constant. Concretely, at each rebalancing date t the optimizer receives $\hat{\mu}$ estimated as the sample mean of daily excess returns $r_{i,\tau} - r_{f,\tau}/252$ computed over the in-sample window $[t-M, t-1]$; the annualized risk-free rate $r_{f,t}$ therefore enters the numerator implicitly through the excess-return series rather than as a separate subtraction inside the optimization routine.

The risk-free rate is therefore a critical and explicitly time-varying input for this optimisation. While some studies, such as Petukhina et al. (2021), assume a constant risk-free rate of zero for simplicity, this analysis deliberately departs from that convention to better reflect the dynamic nature of market conditions. For each monthly rebalancing point in the backtest, $r_{f,t}$ is set equal to the prevailing annualised yield on the 3-Month U.S. Treasury Bill at $t-1$. This ensures that the optimisation at each step uses a timely and realistic measure of the risk-free alternative available to an investor, and that the same time-varying r_f series is consistently reused in the Sharpe-ratio and excess-return computations of Section 4.5.

4.3.3 Utility-maximising portfolios

To generalise the Markowitz framework, we can model the optimal portfolio for an investor with a specific risk preference by maximising a classic mean-variance utility function. This approach provides a flexible bridge between the two extremes of the efficient frontier (GMV and MSR).

$$\max_{w} U(w) = w^T \mu - \frac{\gamma}{2} w^T \Sigma w \quad \text{s.t.} \quad w^T \mathbf{1} = 1,\ w_i \geq 0$$

where γ is the coefficient of absolute risk aversion. The choice of γ is subjective, reflecting an investor's tolerance for risk. A higher γ indicates a greater penalty for variance, pushing the portfolio closer to the GMV, while a lower γ prioritises expected return, moving it towards the MSR. Its advantage is its customisability to an investor's profile, but it inherits the same critical flaw as the MSR portfolio: a strong sensitivity to errors in the expected return vector μ.

To analyse the strategy's behaviour under different investor profiles, the optimisation is run with *multiple* values of γ, each corresponding to a distinct risk-aversion regime, rather than committing to a single representative coefficient.[1] The three resulting investor types are characterised as follows:

- **Return-seeking investor ($\gamma = 0.5$):** Exhibits a low risk aversion, willing to accept higher volatility in pursuit of higher expected returns.
- **Moderately risk-averse investor ($\gamma = 2$):** Represents a balanced, baseline profile, consistent with related academic literature.
- **Highly risk-averse investor ($\gamma = 5$):** Places a strong emphasis on minimising volatility, pushing the allocation closer to the GMV solution.

By evaluating these distinct profiles, we can gain a deeper understanding of how the utility-maximising strategy adapts its allocations in the cryptocurrency market. The same set of γ values is also reused in the certainty-equivalent (CEQ) computation of Section 4.5, so that performance is reported consistently across investor types.

4.4 Smart beta and risk-based strategies

This section explores modern strategies designed to address the challenges of mean-variance optimisation, particularly the instability of expected return estimates. These methods, part of the "Smart Beta 2.0" framework (Amenc, Goltz, and Martellini 2013), construct portfolios based solely on their risk characteristics, using the estimated covariance matrix or its derivatives, thereby avoiding the notoriously noisy μ vector.

4.4.1 Maximum decorrelation (MDC) portfolio

The Maximum Decorrelation strategy builds the most diversified portfolio possible by focusing exclusively on the correlation structure of the assets. It does so by minimising portfolio volatility under the simplifying and unrealistic assumption that all individual asset volatilities are identical (Christoffersen et al. 2010). The key advantage of this deliberate simplification is that it isolates the impact of correlations, creating a portfolio diversified across assets with low correlation to one another. The primary disadvantage is that it completely ignores individual asset volatilities, potentially overweighting highly volatile assets if they happen

[1] Specifically, the three risk-aversion levels tested in the empirical analysis (Chapter 5) are $\gamma \in \{0.5, 2, 5\}$, corresponding respectively to a return-seeking, a moderate, and a highly risk-averse investor profile.

to be good diversifiers from a correlation perspective. The optimisation problem becomes:

$$\min_w w^T \Omega w \quad \text{s.t.} \quad w^T \mathbf{1} = 1,\ w_i \geq 0$$

where Ω is the asset correlation matrix. This strategy has a closed-form unconstrained solution $w^* \propto \Omega^{-1}\mathbf{1}$, making it computationally efficient (Goltz and Sivasubramanian 2018).

4.4.2 Maximum diversification ratio (MDR) portfolio

The Maximum Diversification portfolio, proposed by Choueifaty and Coignard (2008), aims to maximise the portfolio's Diversification Ratio (DR), defined as the ratio of its weighted average asset volatility to its total portfolio volatility:

$$\text{DR}(w) = \frac{w^T \boldsymbol{\sigma}}{\sqrt{w^T \Sigma w}}$$

where $\boldsymbol{\sigma}$ is the vector of individual asset volatilities. Maximising this ratio results in a portfolio that most effectively reduces risk through diversification. The central strength of this approach is its clear objective: to maximize the efficiency of diversification by securing the highest possible benefit relative to the resources allocated. Unlike GMV, which can concentrate heavily in a few low-volatility assets, MDR often produces more balanced portfolios. However, it can also be sensitive to estimation errors in individual volatilities and correlations and may tilt towards higher-volatility assets that offer strong diversification benefits. The optimisation problem is equivalent to maximising the DR:

$$\max_w \frac{w^T \boldsymbol{\sigma}}{\sqrt{w^T \Sigma w}} \quad \text{s.t.} \quad w^T \mathbf{1} = 1,\ w_i \geq 0$$

4.4.3 Risk parity and risk budgeting portfolios

The concept of risk budgeting offers a powerful alternative to capital allocation by focusing instead on the allocation of risk. The central idea is to construct a portfolio where each asset contributes in a balanced way to the total portfolio risk, avoiding the concentration risk found in VW portfolios where most risk often comes from a few dominant assets.

The risk contribution (RC) of an asset i to the total portfolio risk (volatility σ_p) is defined as its weight multiplied by its marginal contribution to risk:

$$\text{RC}_i = w_i \times \frac{\partial \sigma_p}{\partial w_i} = w_i \frac{(\Sigma w)_i}{\sigma_p}$$

where $\sigma_p = \sqrt{w^T \Sigma w}$. For risk measures like volatility, the sum of individual risk contributions equals the total portfolio risk, a property known as Euler decomposition.

4.4.3.1 Equal risk contribution (ERC) portfolio

The Equal Risk Contribution portfolio is the most prominent implementation of the risk parity principle. It is a special case of risk budgeting where the risk budget is allocated equally to each of the N assets. The portfolio is thus the unique weight vector w that solves the condition $\mathrm{RC}_i = \mathrm{RC}_j$ for all assets i, j, subject to the portfolio constraints. Its primary advantage is creating a highly balanced portfolio in terms of risk, ensuring that no single asset dominates the portfolio's volatility profile. This makes it more stable over time than traditional MVO portfolios. Its main challenges are that it is computationally intensive, as it does not have a closed-form solution and must be solved using a numerical optimiser (Maillard, Roncalli, and Teïletche 2010), and it tends to favour low-volatility assets, which may lead to underperformance in strong bull markets. In this work, the ERC weights are obtained numerically by minimising the squared deviations between pairwise risk contributions via the Sequential Least Squares Programming (SLSQP) algorithm (Kraft 1988), as implemented in `scipy.optimize.minimize` (The SciPy community 2026), under the same long-only and full-investment constraints used by the other strategies.

4.4.3.2 Diversified risk parity (DRP) portfolio

The DRP is a simplified heuristic of the ERC portfolio, derived under the strong (and unrealistic) assumption that all pairwise correlations are identical (Amenc, Goltz, and Martellini 2013). This simplification makes the problem trivial to solve, producing a portfolio weighted solely by the inverse of individual asset volatilities. The major advantage of this approach is its stability; by ignoring the complex and often noisy correlation structure, it drastically reduces the potential for estimation error. The clear disadvantage is that it discards valuable information about how assets move together, which is the very essence of diversification. The optimal weights are calculated directly as:

$$w_i = \frac{1/\sigma_i}{\sum_{j=1}^{N} 1/\sigma_j}$$

4.5 Evaluation metrics

The performance of all constructed portfolios will be assessed using a set of metrics that capture return, risk, diversification, and practical trading considerations, drawing on frameworks used in the literature, for instance by Caporin et al. (2014) and Petukhina et al. (2021).

4.5.1 Performance and risk-adjusted return

4.5.1.1 Sharpe ratio (SR)

The Sharpe ratio (Sharpe 1964) is the most widely used measure of risk-adjusted return, quantifying the excess return earned per unit of total risk (volatility).

$$\text{SR} = \frac{\mu_p - r_f}{\sigma_p}$$

4.5.1.2 Adjusted sharpe ratio (ASR)

Standard financial returns are often not normally distributed, exhibiting skewness and fat tails (kurtosis). The Adjusted Sharpe Ratio (Pezier and White 2008) explicitly incorporates penalties for these higher moments, providing a more accurate measure of risk-adjusted performance for strategies with significant tail risk.

$$\text{ASR} = \text{SR}\left[1 + \left(\frac{S}{6}\right)\text{SR} - \left(\frac{K-3}{24}\right)\text{SR}^2\right]$$

where S is the skewness and K is the kurtosis of the portfolio's return series.

4.5.1.3 Certainty equivalent (CEQ)

The certainty equivalent return translates a portfolio's risk and return into a single utility score for a given investor. It represents the guaranteed rate of return an investor would accept rather than taking on the risk of the investment. It provides a direct way to compare portfolios based on an investor's specific risk aversion γ.

$$\text{CEQ} = \mu_p - \frac{\gamma}{2}\sigma_p^2$$

4.5.2 Diversification metrics

To assess how well the portfolios spread exposure and mitigate concentration risk, we employ two distinct diversification metrics. These measures evaluate diversification from the perspective of capital allocation (weight concentration) and risk allocation (correlation benefits).

4.5.2.1 Effective number of assets (Effective N)

This metric measures the concentration of a portfolio's capital allocation. It is calculated as the inverse of the Herfindahl-Hirschman Index of the portfolio weights. A portfolio of N assets that are equally weighted has an Effective N of N, its maximum possible value, while a portfolio fully concentrated in a single asset has an Effective N of 1. A higher value indicates a less concentrated, more diversified portfolio in terms of weight distribution. This definition is consistent with the measure of portfolio concentration used in Amenc, Goltz, and Martellini (2013).

$$N_{\text{eff}} = \left(\sum_{i=1}^{N} w_i^2 \right)^{-1}$$

4.5.2.2 Portfolio diversification index (PDI)

The Portfolio Diversification Index (PDI), introduced by Rudin and Morgan (2006), measures the effective number of independent risk factors in the eligible asset universe at each rebalancing date, computed from the principal components of the $N \times N$ asset covariance matrix Σ estimated over the look-back window. Let Σ be the covariance matrix of the eligible assets and let $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_N$ denote its eigenvalues sorted in descending order. Define the share of variance explained by the k-th principal component as

$$W_k = \frac{\lambda_k}{\sum_{j=1}^{N} \lambda_j}, \qquad \sum_{k=1}^{N} W_k = 1.$$

The PDI is then defined as

$$\text{PDI} = 2 \sum_{k=1}^{N} k \, W_k - 1,$$

and is bounded $1 \leq \text{PDI} \leq N$. Intuitively, PDI captures how evenly portfolio variance is distributed across orthogonal risk factors: it equals 1 when a single principal component absorbs all of the variance (no diversification) and reaches N when every component contributes equally (maximal diversification). The full derivation and the principal-component justification of the formula are given in Rudin and Morgan (2006); here PDI is used purely as a complementary diagnostic to N_{eff}, since the latter measures concentration in the *capital* dimension while PDI measures concentration in the *risk* dimension. The Maximum Diversification Ratio (MDR) strategy of Choueifaty and Coignard (2008), by contrast, is constructed to maximise the closely related *diversification ratio* $w^T\boldsymbol{\sigma}/\sqrt{w^T\Sigma w}$ introduced earlier in this chapter, and is reported alongside PDI rather than identified with it.

4.5.3 Portfolio stability and cost metrics

4.5.3.1 Turnover (TO)

Portfolio turnover is a critical measure of implementability, as higher turnover implies higher transaction costs (fees, slippage) which can erode returns. Following Petukhina et al. (2021), it measures the total fraction of the portfolio that is traded to achieve the target rebalance. It is calculated as:

$$\text{TO}_t = \sum_{i=1}^{N} |w_{i,t} - w_{i,t-1}^{+}|$$

where $w_{i,t}$ is the target weight for asset i at time t and $w_{i,t-1}^{+}$ is the effective weight of asset i just before rebalancing, after accounting for market price changes during the previous period.

4.5.3.2 Target turnover (TTO)

Target turnover measures the instability of the allocation strategy itself, independent of market movements. It quantifies the change in the optimal target weights from one period to the next. A high TTO indicates that the strategy is sensitive to small changes in the input data, which is often a sign of poor stability.

$$\text{TTO}_t = \sum_{i=1}^{N} |w_{i,t} - w_{i,t-1}|$$

References

Amenc, Noël, Felix Goltz, and Lionel Martellini. 2013. *Smart Beta 2.0*. EDHEC-Risk Institute.

Caporin, Massimiliano, Gauthier M. Jannin, Francesco Lisi, and Bertrand B. Maillet. 2014. "A Survey on the Four Families of Performance Measures." *Journal of Economic Surveys* 28(5): 917–942.

Chopra, Vijay K., and William T. Ziemba. 1993. "The Effect of Errors in Means, Variances, and Covariances on Optimal Portfolio Choice." *The Journal of Portfolio*

Management 19(2): 6–11.

Choueifaty, Yves, and Yves Coignard. 2008. "Toward Maximum Diversification." *The Journal of Portfolio Management* 35(1): 40–51.

Christoffersen, Peter, Vihang Errunza, Kris Jacobs, and Xisong Jin. 2010. "Is the Potential for International Diversification Disappearing?" *SSRN Electronic Journal*.

DeMiguel, Victor, Lorenzo Garlappi, and Raman Uppal. 2009. "Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy?" *The Review of Financial Studies* 22(5): 1915–1953.

Goltz, Felix, and Sivagaminathan Sivasubramanian. 2018. *Scientific Beta Maximum Decorrelation Indices*. ERI Scientific Beta.

Kraft, Dieter. 1988. *A Software Package for Sequential Quadratic Programming*. Technical Report DFVLR-FB 88-28. Oberpfaffenhofen, Germany: DFVLR.

Maillard, Sébastien, Thierry Roncalli, and Jérôme Teïletche. 2010. "The Properties of Equally Weighted Risk Contribution Portfolios." *The Journal of Portfolio Management* 36(4): 60–70.

Markowitz, Harry. 1952. "Portfolio Selection." *The Journal of Finance* 7(1): 77–91.

Michaud, Richard O. 1989. "The Markowitz Optimization Enigma: Is 'Optimized' Optimal?" *Financial Analysts Journal* 45(1): 31–42.

Petukhina, Alla, Simon Trimborn, Wolfgang Karl Härdle, and Hermann Elendner. 2021. "Investing with Cryptocurrencies – Evaluating Their Potential for Portfolio Allocation Strategies." *Quantitative Finance* 21(11): 1825–1853. https://doi.org/10.1080/14697688.2021.1880023

Pezier, Jacques, and Anthony White. 2008. "The Relative Merits of Investable Hedge Fund Indices and of Funds of Hedge Funds in Optimal Passive Portfolios." *The Journal of Alternative Investments* 10(4): 37–49.

Rudin, Alexander M., and Jonathan S. Morgan. 2006. "A Portfolio Diversification Index." *The Journal of Portfolio Management* 32(2): 81–89. https://doi.org/10.3905/jpm.2006.611807

Sharpe, William F. 1964. "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk." *The Journal of Finance* 19(3): 425–442.

The SciPy Community. 2026. *`scipy.optimize.minimize`: `SLSQP` Method*. Accessed 30 April 2026. https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.minimize.html

5 Empirical analysis and discussion

5.1 Introduction

This chapter moves from the methodological framework in Chapter 4 and the data construction rules in Chapter 3 to an empirical evaluation of the portfolio strategies over the 2015–2023 period. All strategies follow the same implementation protocol (monthly rebalancing, no short-selling) and are assessed with the metric set defined in Section 4.5: Sharpe ratio (SR), Adjusted Sharpe ratio (ASR), and Certainty Equivalent (CEQ) for risk-adjusted performance; the effective number of assets (Effective N) and the portfolio diversification index (PDI) for diversification; and turnover (TO) and target turnover (TTO) for implementability. The results are therefore read as a joint balance between risk-adjusted performance and practical stability, rather than as a single ranking by return.

The risk-free rate used in the tangency portfolio and in excess-return calculations is the 3-Month U.S. Treasury Bill series. The series is aligned to the daily index and forward-filled across non-trading days to ensure that the crypto and ETF calendars remain comparable. Forward-filling is the standard approach in financial research when dealing with asynchronous trading calendars: the risk-free rate published on a given day reflects the prevailing short-term funding cost and remains economically valid until the next observation. Because 3-month T-bill rates change slowly relative to daily equity or crypto returns, the interpolation error is negligible compared to alternative methods such as linear interpolation, which would introduce artificial smoothness, or omission, which would bias excess-return calculations on weekends and holidays (Campbell, Lo, and MacKinlay 1997). When computing daily excess returns in datasets where assets trade on different calendars (e.g., ETFs that follow exchange holidays versus crypto that trades continuously), it is necessary to align the risk-free benchmark to the return series. The standard practical convention is to treat the most recently observed short rate as valid until the next observation becomes available, i.e., to carry forward the last available value rather than interpolating or leaving missing days. This preserves a consistent definition of excess return at the chosen sampling frequency and avoids introducing artificial dynamics into the risk-free series. This approach is consistent with the general econometric treatment of return measurement and time-series construction in empirical finance (Campbell, Lo, and MacKinlay 1997).

To isolate the effect of the asset universe, the backtest is conducted across three scenarios:

1. **Scenario A: Crypto-only (dynamic).** The eligible set changes monthly according to the survivorship and 1% relative market capitalization filters in Section 3.5.1.

2. **Scenario B: Traditional-only (static).** A fixed set of seven highly liquid instruments (Section 3.3) is used as a proxy for core asset classes.

3. **Scenario C: Mixed (hybrid).** The dynamic crypto sleeve is combined with the static traditional basket to test the diversification effects of blending the two regimes.

The main empirical claim is *regime-dependent optimality*: the strategy that maximizes risk-adjusted performance depends on whether the universe is dynamic and heavy-tailed (crypto), static and more stable (traditional), or a hybrid of both. This concept recognizes that no single allocation rule dominates across all market environments. In environments characterized by extreme tail risk and frequent constituent changes (crypto), strategies that minimize variance outperform because estimation error in expected returns is amplified by heavy tails. In stable, mature markets (traditional assets), strategies that capture systematic risk premia (such as equity beta) dominate because input estimates are more reliable.

In hybrid environments, strategies that exploit correlation structures emerge as optimal because they balance diversification benefits against the instability of the dynamic sleeve. Each subsection explicitly links the narrative to the metrics in Section 4.5, so that the comparisons are grounded in the same evaluation language across regimes.

All results in this chapter are reported *gross of transaction costs*: trading fees, bid–ask spreads and slippage are explicitly excluded from the returns and from all derived metrics, so that the comparison across strategies isolates the statistical properties of each allocation rule rather than the incremental cost of implementing it.[1] Turnover (TO) and target turnover (TTO) are nevertheless reported throughout this chapter as proxies for the cost intensity of each strategy, so that the cost-free results can already be read with the implementability of each rule in mind.

Before proceeding to the scenario analysis, it is useful to clarify the interpretation of the core metrics. The *Sharpe ratio (SR)* measures annualized excess return per unit of volatility and serves as the baseline risk-adjusted measure; values above 0.5 are generally considered acceptable, while values above 1.0 indicate strong risk-adjusted performance. The *Adjusted Sharpe ratio (ASR)* penalizes negative skewness and excess kurtosis, providing a more conservative estimate in fat-tailed distributions; a large gap between SR and ASR signals that the strategy's returns are driven by infrequent large gains that may not persist. The *Certainty Equivalent (CEQ)* expresses the annualized return that a mean-variance investor would accept with certainty in lieu of the risky portfolio; negative CEQ indicates that the strategy's volatility is too high relative to its return for the assumed risk aversion level. The *Effective N* captures the diversification implied by the weight distribution, where values close to 1 indicate concentration in a single asset and higher values indicate broader dispersion. The *Portfolio Diversification Index (PDI)* measures the ratio of weighted average asset volatility to portfolio volatility, with values above 1 indicating diversification benefit. Finally, *Turnover (TO)* and *Target Turnover (TTO)* quantify trading intensity: TO accounts for market drift between rebalances, while TTO measures pure allocation changes; lower values imply easier implementation.

[1] This is a known and deliberate limitation of the Chapter 5 analysis. Realistic execution costs (Binance fees on the crypto sleeve, IBKR commissions on the ETF sleeve and a slippage proxy) are introduced and quantified in the operational backtest of Chapter 6, where the same strategies are re-evaluated net of frictions and the cost-induced gap to the Chapter 5 paper-portfolio results is decomposed explicitly.

5.2 Scenario A: The crypto-only universe (dynamic)

The distributional evidence for the eligible cryptocurrency assets is reported in Chapter 3, especially in Table 3.7. Scenario A is the most unstable regime because the investable universe rotates monthly through the 1% relative market-capitalisation filter in section 3.5. This structural churn is the first-order source of estimation error: each rebalance combines heavy-tailed returns with a changing asset set. The setting therefore provides the sharpest test of Michaud (1989): return-seeking optimisation should amplify noisy sample means rather than reveal persistent alpha.

5.2.1 Performance and return dynamics

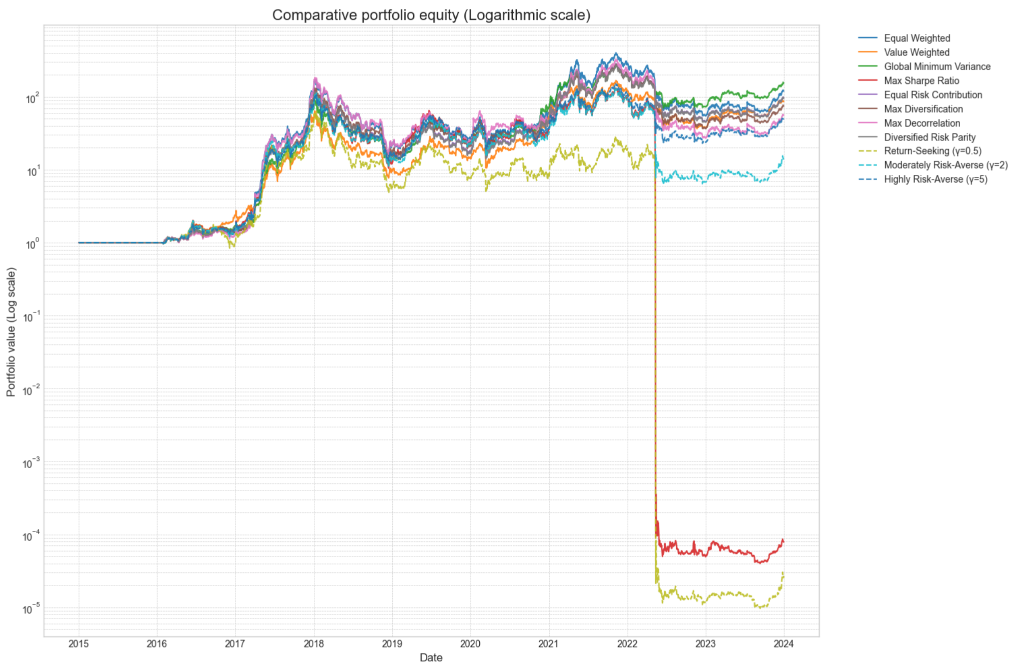

Figure 5.1: Comparative equity curves for the crypto-only universe (log scale).

GMV dominates Scenario A on SR, ASR, CEQ and annualised return, confirming that variance minimisation is the appropriate response to the heavy-tailed and rotating crypto universe documented in Chapter 3. VW ranks second on most performance metrics because concentration in the largest tokens, especially the Bitcoin and Ether complex, provides scale exposure without estimating expected returns. MSR and the Return-Seeking utility profile generate episodic spikes in the equity curves, but their negative ASR and CEQ show that these gains are more than offset by skewness, kurtosis and drawdown risk. The SR–ASR

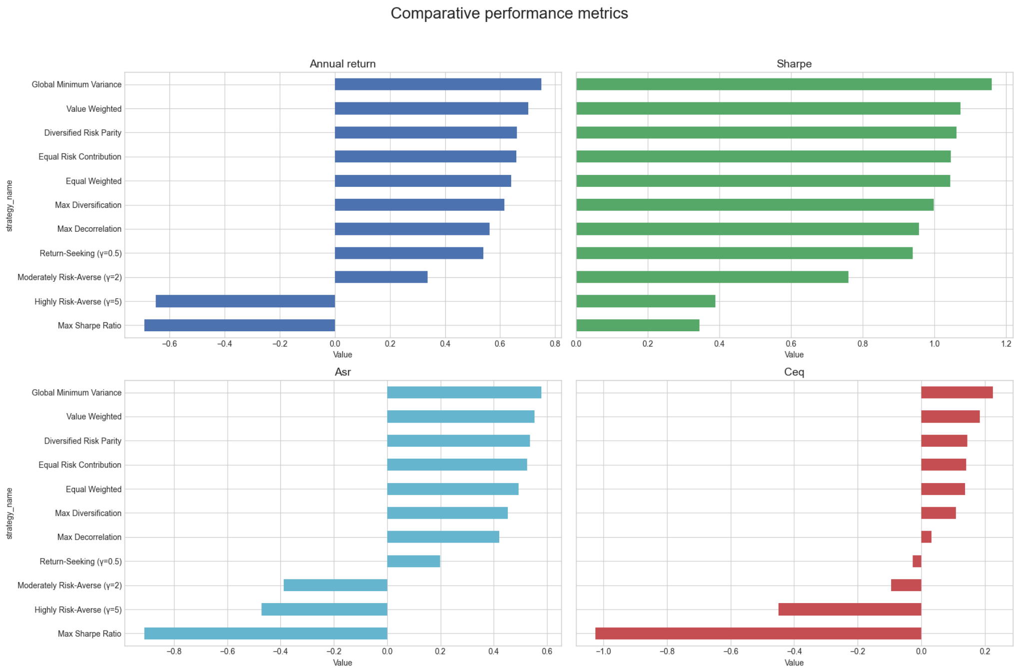

Figure 5.2: Performance metrics for the crypto-only universe.

gap is therefore the empirical form of the Michaud (1989) error-maximisation problem: noisy return estimates are amplified into unstable portfolios whose out-of-sample distribution is less symmetric and less persistent than SR alone suggests. The allocation mechanics driving this hierarchy are visible in the composition heatmap.

5.2.2 Portfolio composition

GMV and the risk-averse utility profiles concentrate in Bitcoin, which acts as the lowest-volatility anchor in the crypto-only universe and usually absorbs a large share of the average target weight. MSR and Return-Seeking spread weight across higher-beta altcoins; their Eff-N is therefore higher, but the gain is capital dispersion rather than risk-adjusted diversification. This distinction is central to Scenario A: a broader weight vector does not automatically improve SR, ASR or CEQ when each additional asset brings unstable tail risk. VW offers an intermediate mechanism by preserving exposure to dominant tokens without running a return forecast. The cost of this dispersion is visible in the turnover statistics.

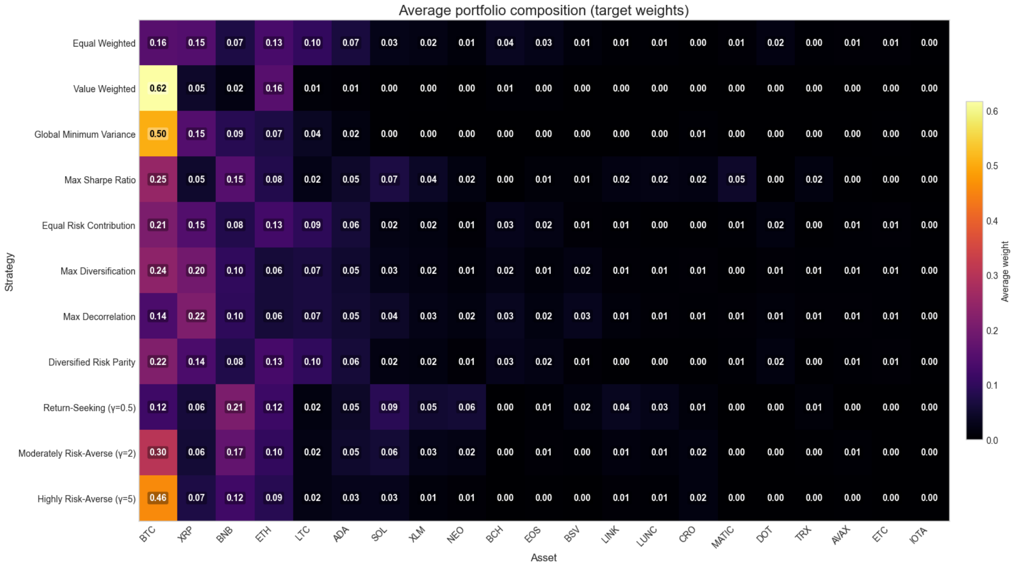

Figure 5.3: Average portfolio composition in the crypto-only universe.

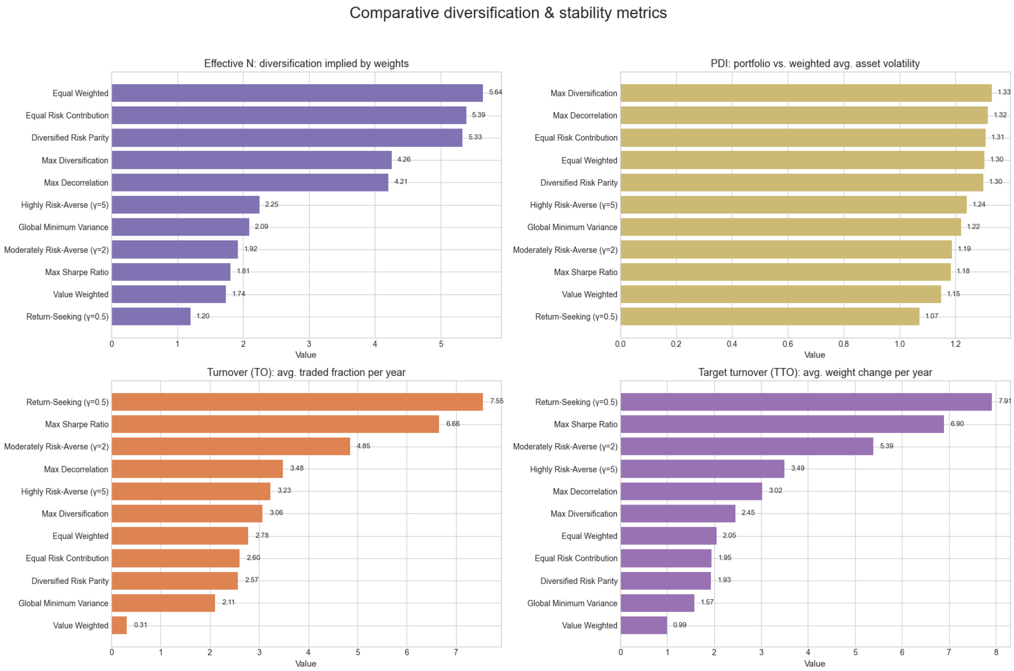

Figure 5.4: Diversification and stability metrics for the crypto-only universe.

5.2.3 Stability and turnover

MSR and Return-Seeking have the highest TO and TTO in Scenario A, so the most aggressive objectives are also the least implementable before transaction costs are even applied. GMV has the best alignment between strong SR, ASR and CEQ and lower trading intensity, which makes its dominance both statistical and operational. PDI remains non-trivially above one for GMV and ERC, confirming that diversification benefit can exist even when the portfolio is concentrated in Bitcoin, because Bitcoin is not perfectly correlated with the altcoin set. The gap between TO and TTO also signals that market drift compounds the instability created by the dynamic eligibility rule. In Scenario B, the investable universe is static and distributional inputs are more reliable; the performance hierarchy is reversed.

5.3 Scenario B: The traditional-only universe (static)

The descriptive properties of the traditional instruments are documented in Chapter 3, including the seven-instrument universe in Table 3.4. Scenario B is the direct contrast to Scenario A: the universe is static, tail risk is lower, and input estimates are therefore more reliable. This is the environment in which mean-variance theory is expected to be least fragile. The dominant outcome is VW, because the traditional basket rewards persistent exposure to the equity risk premium rather than defensive variance minimisation.

5.3.1 Performance and return dynamics

VW dominates Scenario B on SR, ASR and CEQ and delivers the most consistent equity path, because SPY carries the equity risk premium over the 2015–2023 horizon while bonds and commodities act as secondary buffers. GMV sacrifices return by over-allocating to GOVT and USIG, so its low-volatility objective becomes an inefficient flight to safety when tail risk is already limited. The SR–ASR gap is much smaller than in Scenario A for every strategy, indicating that the traditional universe is closer to the distributional environment in which the Section 4.5 performance metrics are mutually consistent. MSR still does not become dominant, because small perturbations in estimated expected returns continue to affect its weights. The composition heatmap clarifies the mechanism.

5.3.2 Portfolio composition

VW allocates primarily to SPY, explaining its superior SR and CEQ as a market-cap-based exposure to the equity premium with bond, commodity and gold

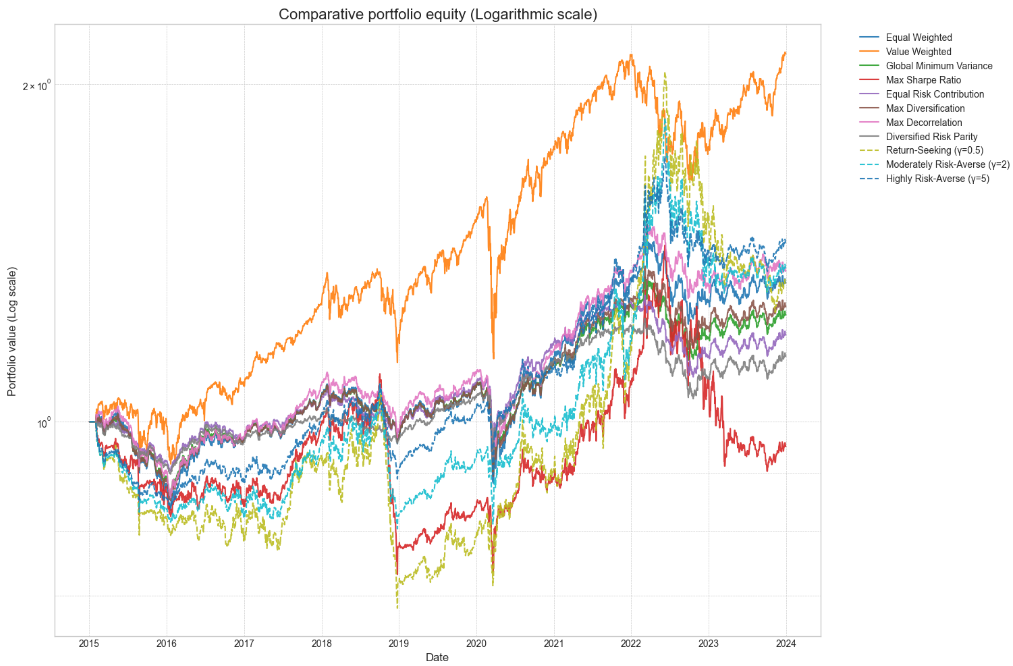

Figure 5.5: Comparative equity curves for the traditional-only universe (log scale).

Figure 5.6: Performance metrics for the traditional-only universe.

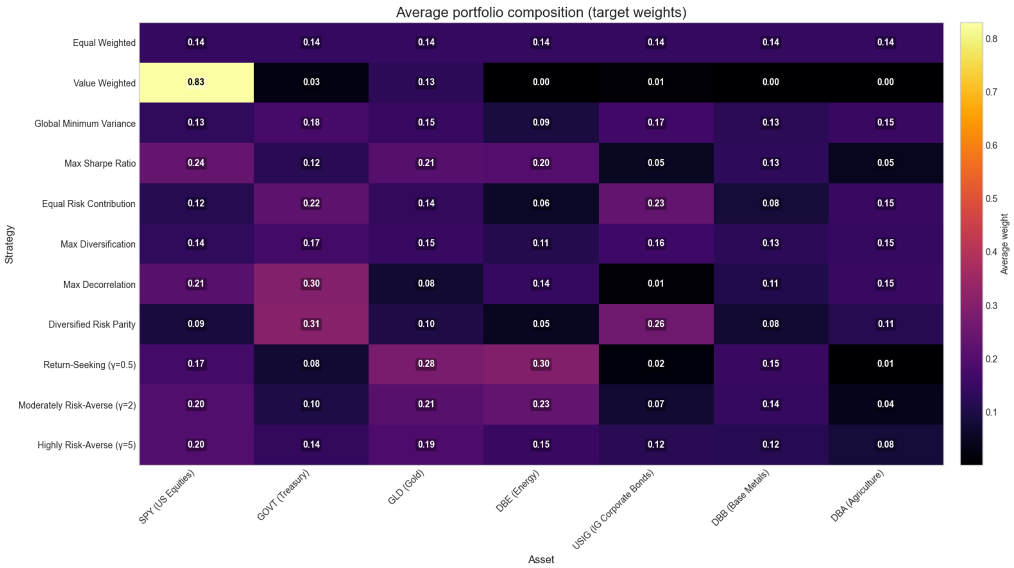

Figure 5.7: Average portfolio composition in the traditional-only universe.

buffers. GMV and the risk-averse utility portfolios allocate heavily to GOVT and USIG, which lowers volatility but gives up too much return in a regime where extreme tail risk is not the binding constraint. ERC and MDC achieve higher Eff-N and PDI through broader dispersion across asset classes, yet this diversification is not rewarded by the performance metrics. Scenario B therefore reverses the crypto result: when estimates are more stable, the opportunity cost of leaving equity beta is larger than the marginal diversification gain. Turnover confirms the stability advantage of the static universe.

5.3.3 Stability and turnover

TO and TTO are materially lower in Scenario B than in Scenario A for every strategy, confirming that the absence of constituent rotation removes the main structural source of trading intensity. Even MSR, which is the most unstable return-seeking rule in the crypto-only case, becomes more manageable when the asset set is fixed and estimates are less noisy. The remaining turnover is therefore driven primarily by re-optimisation rather than by eligibility churn. This distinction is important for Chapter 6, where TO is translated into execution costs: Scenario A's implementability problem is not only an optimiser problem but also a universe-design problem. Scenario C combines both universes and tests whether the diversification potential of the crypto sleeve can be exploited without inheriting its full instability.

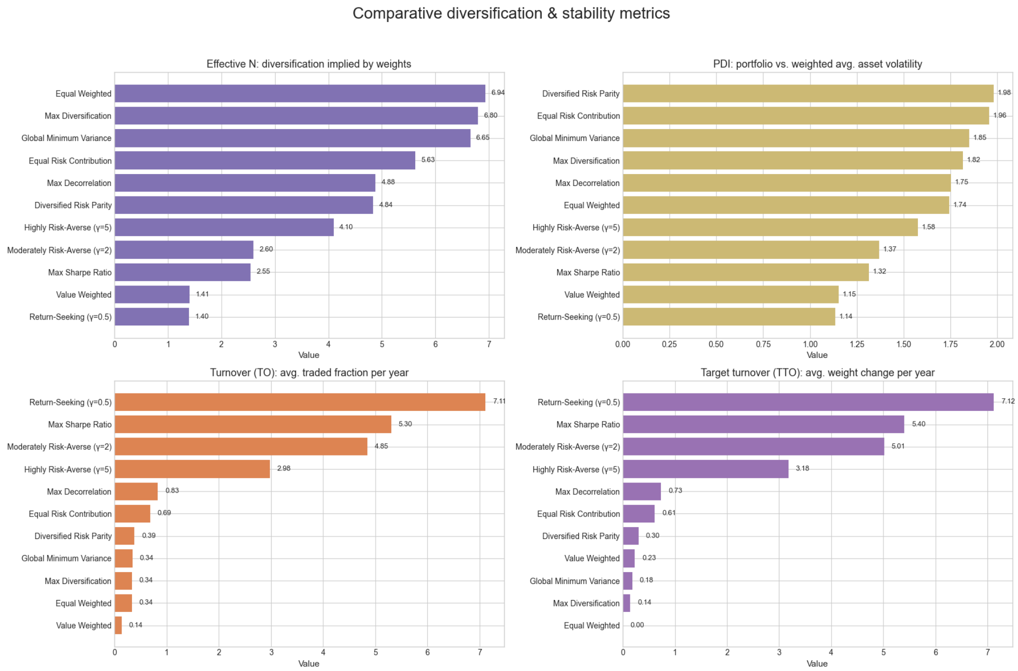

Figure 5.8: Diversification and stability metrics for the traditional-only universe.

5.4 Scenario C: The mixed-asset universe (hybrid)

Chapter 3 documents both the seven traditional instruments in Table 3.4 and the dynamic crypto sleeve generated by section 3.5. Scenario C pools these assets in a single optimisation: all eligible traditional and crypto assets enter the same covariance matrix, with no sleeve cap or two-stage hierarchy. This design lets cross-sleeve correlations determine the weights directly (Brière, Oosterlinck, and Szafarz 2015). MDC is the dominant strategy because, in a mixed regime, correlation structure matters more than marginal expected returns (Choueifaty and Coignard 2008).

5.4.1 Performance and return dynamics

In Scenario C, no single strategy dominates all metrics simultaneously. MDR and ERC rank above MDC on SR, but their ASR is compressed by the heavier tail exposure that comes with broader crypto allocation. MDC achieves the most favourable balance: among strategies with meaningful crypto sleeve engagement, it delivers the highest CEQ and keeps TO and TTO within implementable bounds, while avoiding the boom-and-bust profile of return-seeking allocations. Return-Seeking and MSR show the widest equity swings: their crypto exposure produces episodic gains, but volatility and higher moments compress ASR and

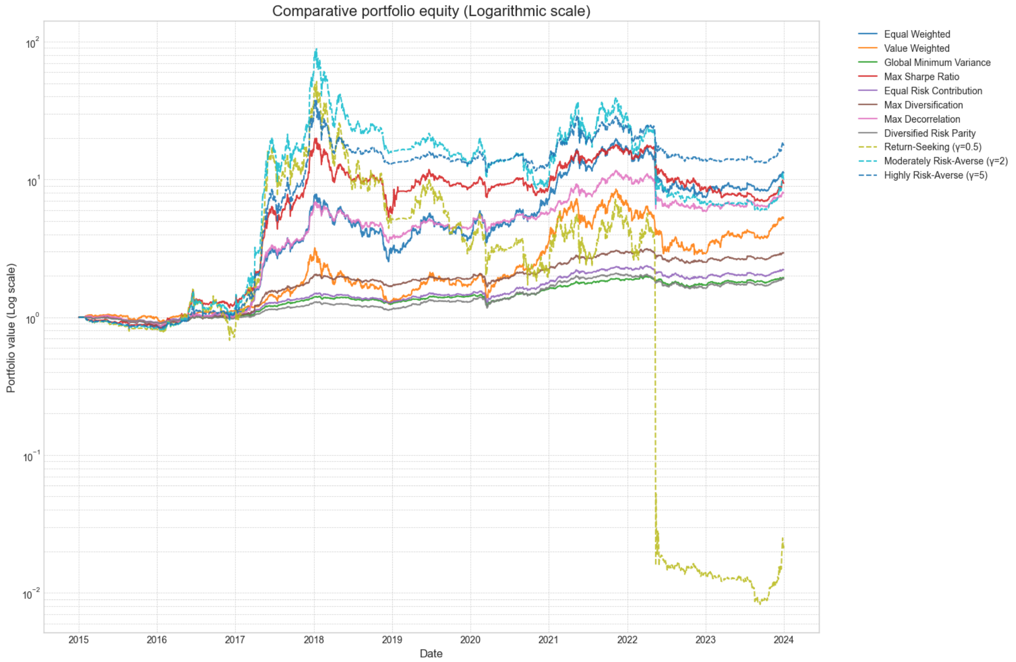

Figure 5.9: Comparative equity curves for the mixed-asset universe (log scale).

Figure 5.10: Performance metrics for the mixed-asset universe.

push CEQ toward an unattractive level. GMV underweights the crypto sleeve because variance dominates its objective, so it behaves more like a defensive traditional portfolio than a mixed allocation. The SR–ASR gap is narrowest for ERC and DRP in Scenario C, while MDC's ASR is penalised by the higher-moment exposure from its crypto sleeve, a cost that does not appear in the SR but is captured in the CEQ comparison. The composition heatmap in the next subsection confirms that MDC achieves this by maintaining a bounded but meaningful crypto allocation.

5.4.2 Portfolio composition

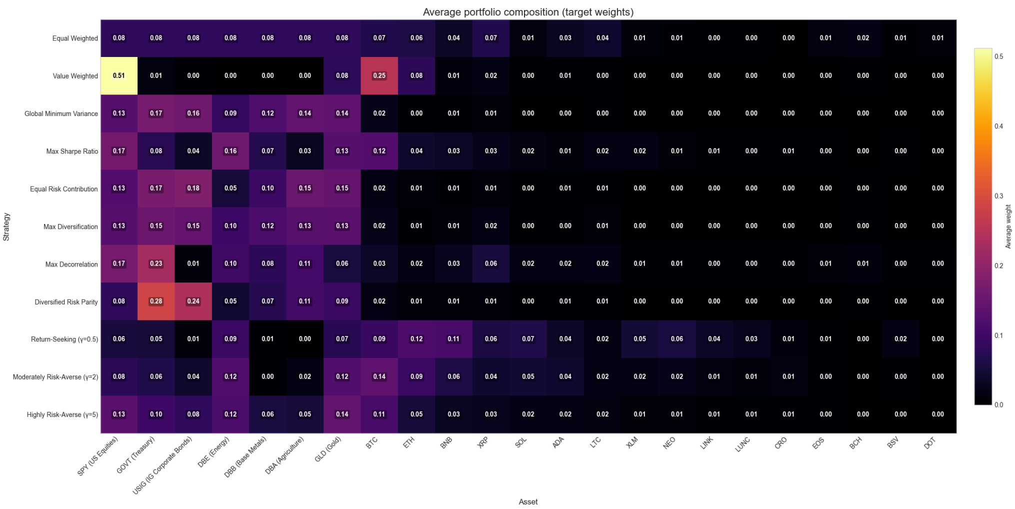

Figure 5.11: Average portfolio composition in the mixed-asset universe.

MDC maintains a genuine cross-sleeve blend rather than a traditional portfolio with a token crypto allocation. Its average weights combine commodity and crypto exposures with traditional anchors, so the allocation is endogenously driven by correlation structure rather than by a fixed sleeve budget. Risk-averse utility portfolios resemble Scenario B because bonds and gold dominate when γ is high, effectively suppressing the crypto sleeve. Return-Seeking and MSR reproduce the Scenario A failure mode by concentrating in high-beta crypto assets. Moderate crypto exposure improves PDI without destabilising TO and TTO, while heavier allocations convert diversification potential into tail exposure. The turnover evidence in the next subsection quantifies the cost of exceeding this bounded-diversification threshold.

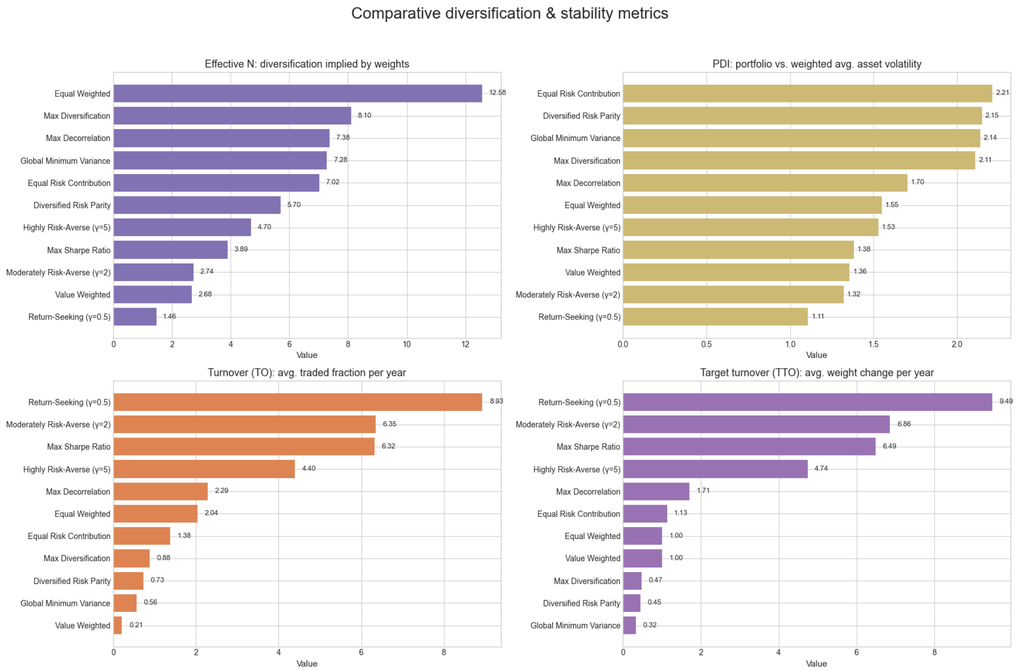

Figure 5.12: Diversification and stability metrics for the mixed-asset universe.

5.4.3 Stability and turnover

TO and TTO are higher than in Scenario B for all strategies that engage the crypto sleeve, but the increase is uneven across allocation philosophies. Correlation-aware strategies, especially MDC and ERC, remain materially more stable than MSR and Return-Seeking, whose turnover inherits the full instability of Scenario A. MDC delivers the most favourable balance in terms of CEQ and implementability: among strategies that engage the crypto sleeve, it avoids the extreme turnover of MSR and Return-Seeking while producing a risk-adjusted return profile that Chapter 6 can plausibly cost. Eff-N and PDI are also strongest in the hybrid universe for ERC and MDC, confirming that the mixed opportunity set contains the greatest diversification potential. The cross-scenario comparison in section 5.5 consolidates these three regime findings and isolates the strategies that drive the regime-dependent narrative.

5.5 Cross-scenario comparison

GMV, VW and MDC represent three distinct allocation philosophies and three distinct compounding profiles. The log scale makes total growth and drawdown depth comparable in the same panel: GMV protects capital best in the crypto-only stress regime, VW compounds most efficiently in the static tradi-

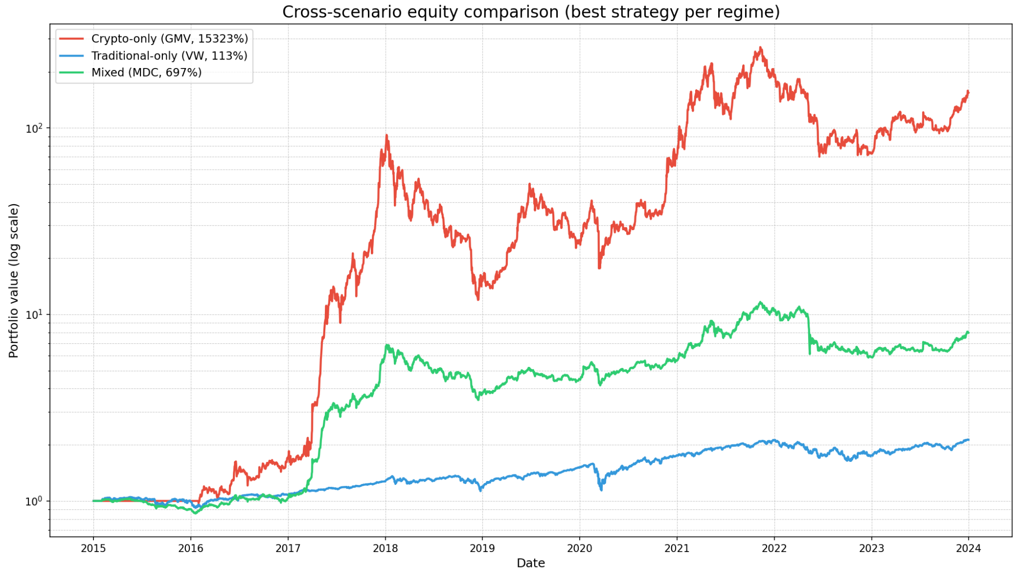

Figure 5.13: Cross-scenario equity comparison: best strategy per regime (log scale; legend reports cumulative return over 2015–2023 in parentheses).

tional regime, and MDC delivers the strongest hybrid path by exploiting cross-sleeve decorrelation. No single curve dominates the full horizon. During the 2022 stress episode, GMV and MDC absorb the drawdown better than VW, while VW benefits from the subsequent recovery in traditional equity beta. The cumulative return statistics in the legend anchor the visual comparison, but the relevant conclusion is not a single winner. The overlay is the first cross-scenario confirmation of regime-dependent optimality, which the next figure tests across the common SR, ASR and CEQ metrics.

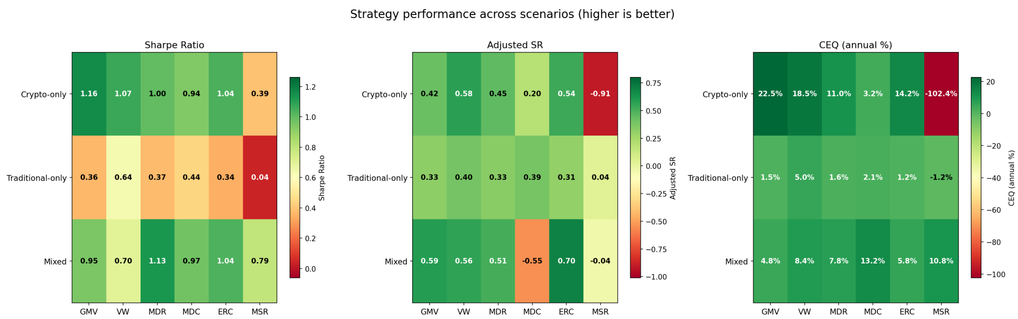

Figure 5.14: Strategy performance comparison across scenarios.

Strategy ranks change sharply across regimes, which rules out any single allocation rule as universally optimal. GMV moves from the top of Scenario A on SR, ASR and CEQ to the bottom of Scenario B on CEQ because its variance objective sacrifices too much equity premium. VW moves in the opposite direction, from weak crypto performance to Scenario B dominance, because a static traditional universe rewards market-cap exposure. MDC is the most consistent middle-to-top performer across scenarios, and in Scenario C it achieves the best CEQ among strategies that actively engage the crypto sleeve, making it the preferred rule from a utility-consistent perspective. MSR never wins a scenario, which directly limits the practical appeal of Sharpe-ratio maximisation under fat-tailed or even mildly unstable inputs. The implementability panel completes the comparison by adding TO, TTO, Eff-N and PDI.

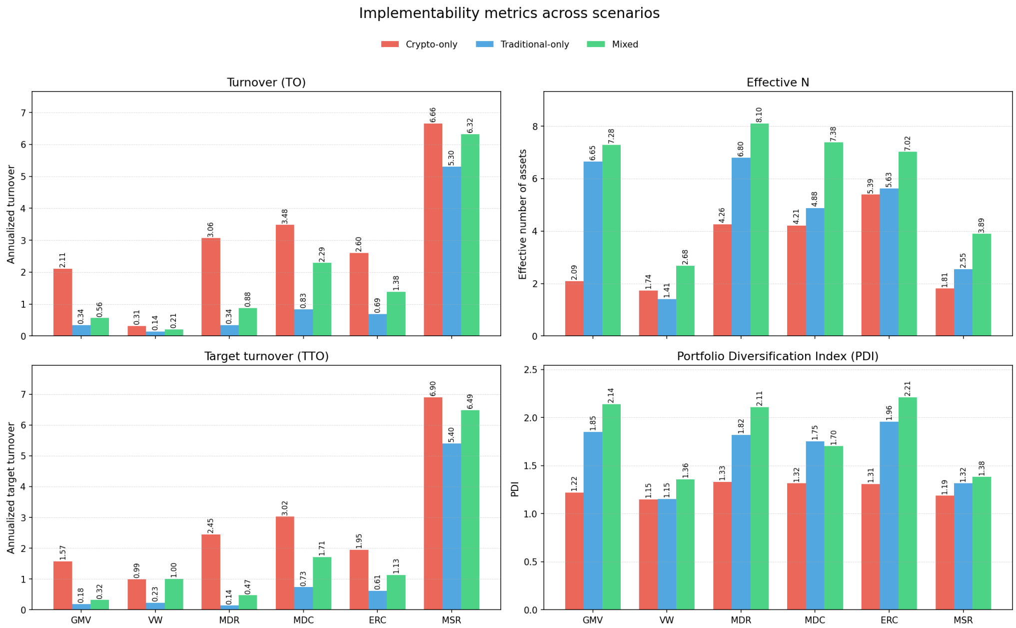

Figure 5.15: Implementability metrics across scenarios: turnover (TO, TTO, left panels) and diversification (Effective$N$, PDI, right panels).

Turnover is structurally highest in Scenario A for every strategy, confirming that constituent rotation is a main driver of cost intensity rather than a minor implementation detail. The gap between TO and TTO is widest for MSR and Return-Seeking, where market drift compounds unstable target allocations. In Scenario B, both metrics fall because the static universe removes constituent-driven churn. Scenario C sits between the two regimes: MDC and ERC increase turnover only moderately relative to Scenario B, while return-seeking rules inherit the instability of the crypto sleeve. The right panels show that Eff-N and PDI are highest for MDC and ERC in the mixed universe, confirming the diversification hypothesis when correlation-aware allocation is used. The optimality summary condenses these trade-offs into a single regime table.

The optimality summary should be read as a compact falsification of universal strategy dominance. Winners are dispersed across scenarios and metrics:

Regime-dependent optimality: best strategy per scenario and metric (3 scenarios x 4 metrics)

Scenario	Best SR	Best ASR	Best CEQ	Lowest TO
Crypto-only	GMV	VW	GMV	VW
Traditional-only	VW	VW	VW	VW
Mixed	MDR	ERC	MDC	VW

Figure 5.16: Regime-dependent optimality summary: best strategy per scenario and metric (3 scenarios $\times$ 4 metrics).

GMV wins when the universe is dynamic and heavy-tailed, VW wins when the universe is static and equity beta is rewarded, and MDC achieves the best CEQ when crypto and traditional sleeves are optimised jointly through correlation management rather than return forecasting. The only stable pattern is the type of rule that matches the regime's statistical structure. This maps directly to the Chapter 4 theory and closes the empirical case for regime-dependent optimality. The same evidence also supports a pragmatic sequencing for mixed institutional sleeves: prioritise stability when the universe is unstable, prioritise market beta when the universe is static, and prioritise correlation management when the universe is mixed, rather than relying on a single universal rule across regimes. Chapter 6 asks how far these rankings survive under transaction costs and execution frictions; that chapter's backtest uses 19 cryptocurrencies instead of 21 because complete Binance Spot daily series were missing for BSV and CRO, as set out in Section 6.2.1.

Event-level detail on the May 2022 Terra/LUNA collapse, including the LUNC-restriction robustness checks, is relegated to Appendix B for readers who want the episode separate from the regime-level narrative.

References

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay. 1997. *The Econometrics of Financial Markets*. Princeton, NJ: Princeton University Press.

Brière, Marie, Kim Oosterlinck, and Ariane Szafarz. 2015. "Virtual Currency, Tangible Return: Portfolio Diversification with Bitcoin." *Journal of Asset Management* 16(6): 365–373. https://doi.org/10.1057/jam.2015.5

Choueifaty, Yves, and Yves Coignard. 2008. "Toward Maximum Diversification." *The Journal of Portfolio Management* 35(1): 40–51.

6 Realistic implementation: transaction costs and robustness

6.1 Introduction

Up to this point, the empirical analysis has focused on a frictionless end-of-day rebalancing protocol. While this setup is a necessary first step because it isolates the statistical effect of the allocation objectives and of the asset-universe design, it is not sufficient to assess economic viability. In practice, portfolio strategies are executed through an order-management and brokerage layer that introduces frictions: explicit fees (commissions), implicit costs (bid–ask spreads and market impact), and implementation constraints (cash drag, minimum order sizes, and execution timing). These frictions are known to be first-order for strategies with high turnover, and they can materially alter both the level and the ranking of risk-adjusted performance (Almgren and Chriss 2001; Chan and Lakonishok 1995). The objective of this chapter is therefore to stress-test the empirical conclusions of Chapter 5 under a more realistic trading model.

The main contribution of the analysis is a *decomposition* of the implementation effect into two components:

- **Pure cost drag**: the mechanical reduction in performance that follows from paying trading costs, holding everything else fixed.
- **Endogenous implementation response**: the change in the realized return path induced by the interaction between costs, discrete execution, and portfolio constraints (e.g., the strategy may trade less, remain partially in cash after rebalances, or reduce exposure after large losses).

This separation matters because observed performance differences under costs need not be monotone in the cost parameter. Even in a long-only setting, adding costs can modify the trading trajectory in ways that partially offset (or, in rare cases, dominate) the direct cost drag. The analysis thus focuses not only on whether costs reduce returns, but on *why* the realized changes occur and how they relate to implementability proxies such as turnover. This framing is aligned with implementation shortfall analysis, which separates explicit and implicit execution costs and emphasizes that portfolio outcomes depend on the execution policy, not only on posted fee schedules (Khandoker, Bhuyan, and Singh 2017).

6.2 Backtest design and cost model

6.2.1 Data

The backtest operates on daily OHLCV data[1]: Open, High, Low, Close prices and traded Volume for each asset in the investment universe. OHLCV is the standard market-data format: the *Open* price is the first trade of the day, the *High* and *Low* are the intraday extremes, the *Close* is the last recorded price, and *Volume* is the total number of units traded during the session.

Cryptocurrency prices are sourced from Binance Spot. For each coin, the loader selects, among USD-stable quote pairs with available files (USDT, USDC, BUSD), the series with the earliest start date, so as to maximize the usable history for that asset; the effective quote currency therefore varies by ticker and is not fixed to USDT. ETF prices are taken from Tiingo (daily closes in USD). Tiingo is used for historical ETF OHLCV because it provides a long, reproducible panel suitable for research backtests; Interactive Brokers (IBKR) is not used as a price vendor here. IBKR enters only as an *external benchmark* for realistic all-in ETF trading costs (Section 6.2.3), not as the source of returns. Mixing Tiingo prices with IBKR-style cost assumptions is standard in reduced-form implementation

[1] OHLCV denotes Open, High, Low, Close prices and Volume; it is the standard market-data format used by trading systems and backtesting engines.

studies; it introduces a minor source-consistency limitation (prices from one vendor, cost calibration from another), which is acceptable for ranking strategies under common frictions but would not support venue-specific microstructure claims.

The sample period runs from January 2015 to December 2023, covering the same time span as Chapters 4 and 5 and the same mixed-universe design (Scenario C in Chapter 5). The data pool comprises 19 cryptocurrencies and 7 ETFs. Two tokens present in the Chapter 5 universe (BSV and CRO) are omitted because no complete Binance Spot daily series was available for them; the remaining 19 cover the same eligibility and filtering logic, and all seven ETFs are retained. Because the dynamic eligibility filter selects at most 12 cryptocurrencies per month, the practical overlap with the Chapter 5 eligible sets is high, and the omission does not materially alter strategy comparisons. The data are stored as individual CSV files per asset, each containing the five OHLCV columns at daily frequency.

6.2.2 Scenario design

The empirical exercise uses the same strategic definitions as in Chapter 4 and the same monthly rebalancing schedule[2] as Chapter 5, but it replaces the purely vectorized portfolio accounting with an order-driven backtesting engine. The strategy set includes EW, ERC, DRP, MDR, MDC, GMV, MSR, and three utility profiles ($\gamma \in \{0.5,\ 2,\ 5\}$). The value-weighted (VW) benchmark used in Chapter 5 is excluded here because it requires daily market-capitalization data, which are not part of the OHLCV feed consumed by the execution engine. Specifically, the experiment runs two scenarios on the same daily OHLCV data:

1. **No-cost baseline**: all-in costs are set to zero; the execution engine still processes discrete orders, so path effects from order sizing and cash rounding are present, but no fees are charged.

2. **Cost-inclusive scenario**: all-in costs are set to realistic values and applied at execution time, as described in Section 6.2.3.

6.2.3 Transaction cost parameterization

Transaction costs are modeled in *basis points* (bps).[3] One basis point equals one hundredth of a percentage point: 1 bp $= 0.01\% = 0.0001$. Costs are applied as a proportional charge on the notional value of each executed order. The "all-in"

[2]Rebalancing frequency is the periodic re-optimization interval. This thesis uses monthly rebalancing to balance turnover control with responsiveness to regime changes.

[3]One basis point is one hundredth of a percentage point; 1 bps $= 0.01\%$.

convention bundles explicit fees (exchange commissions) and estimated implicit execution costs (spread capture, microstructure slippage, and partial market impact) into a single rate per asset class:

- **Cryptocurrencies**: 15 bps (0.15%) all-in. This is calibrated to the spot taker fee schedule published by Binance (10 bps at the base tier), plus a conservative 5 bps allowance for slippage and spread crossing (Binance 2024).

- **ETFs**: 3.5 bps (0.035%) all-in. This reflects a typical institutional all-in execution cost for US-listed ETFs as summarized on Interactive Brokers' commission pages (sub-penny commissions and tight spreads on liquid ETFs) (Interactive Brokers 2024).

This chapter is not meant to provide a venue-by-venue microstructure model; its purpose is to *rank robustness* across strategies under the same execution assumptions. A parsimonious bps model is appropriate for this objective, provided that its economic logic is explicit: strategies with higher trading intensity should be more penalized in net performance, and strategies with stable targets should be less sensitive. These principles are consistent with evidence from institutional trade data, where costs rise with liquidity demand and traded volume (Chan and Lakonishok 1995), and with the optimal execution literature, where trading schedules balance market impact and risk (Almgren and Chriss 2001).

6.2.4 Execution engine: Backtrader

The backtest is implemented using Backtrader,[4] an open-source event-driven backtesting framework for Python (Rodrigues 2024). Unlike the vectorized return arithmetic used in Chapter 5, Backtrader simulates a realistic brokerage environment: a virtual broker holds cash and positions, receives market orders at each rebalancing date, fills them against the daily price bar, and deducts commissions from the cash account at the time of execution.

Concretely, the implementation works as follows. At each monthly rebalancing date, the strategy computes target weights using the same optimization objectives as in Chapter 5 (over a trailing 365-day lookback window). The engine then translates these target weights into market orders for each asset. Orders are executed at the next available price, and the broker deducts the all-in cost (in bps of traded notional, as specified in Section 6.2.3) from the portfolio's cash balance. This process introduces three effects that are absent in a frictionless backtest:

[4] Backtrader is an open-source Python backtesting framework that simulates discrete order execution, broker cash accounting, and commission charging.

- **Cash accounting**: commissions reduce available capital and can induce residual cash holdings after rebalancing, especially under discrete sizing and rounding constraints.
- **Path dependence**: the realized sequence of trades and fills affects future positions and, therefore, future turnover and exposure.
- **Nonlinear interaction with constraints**: under long-only and full-investment constraints, costs can effectively shrink the feasible set, changing the strategy's exposure profile relative to the no-cost ideal.

Monthly rebalancing and why path dependence affects realized returns. The no-cost and cost-inclusive scenarios are *two separate full simulations of the same strategy*, not a single return path with fees subtracted after the fact. On each rebalancing date (aligned with Chapter 5: month-end dates in the eligibility schedule), the strategy computes the same target *weights* from past data. Orders are translated into share quantities using Backtrader's `order_target_percent`, which scales positions to the current portfolio value held by the broker. After commissions are paid in the cost-inclusive run, cash and total value differ from the no-cost run, so the same target percentages imply *different dollar positions and share holdings* from that date onward. Between rebalances, positions evolve only from market moves; the next rebalance again starts from a different cash-and-position state. Over many months, these differences compound, so the realized annualized return in the cost run need not equal the no-cost return minus a simple fee drag. Strategies with high turnover and frequent rebalancing therefore exhibit larger gaps between realized Δr and the *pure* commission approximation. Occasionally the realized cost-inclusive return can exceed the no-cost return in the sample: this is a finite-sample, path-dependent outcome (e.g. different effective exposure after fees), not evidence that trading costs raise expected returns.

These effects are well documented in the market microstructure and optimal execution literature. Empirically, trading costs increase in traded quantity and in liquidity demand, creating a direct link between turnover and net performance (Chan and Lakonishok 1995). Theoretically, proportional transaction costs generate regions of inaction ("no-trade" wedges) in optimal portfolio choice, implying that the presence of costs alters rebalancing frequency and target-tracking behavior (Davis and Norman 1990; Leland 1999). Dynamic multi-asset models with trading frictions reach a similar conclusion: optimal policies internalize adjustment costs and therefore smooth trading relative to frictionless targets (Gârleanu and Pedersen 2016). The purpose of the present analysis is not to solve the optimal control problem, but to evaluate whether the strategies identified as stable in Chapter 5 remain so once these frictions are introduced.

6.3 Measuring the cost impact: realized vs. pure drag

Let r^{nocost} and r^{cost} denote the annualized realized returns produced by the execution engine under the no-cost and cost-inclusive scenarios, respectively. The *realized* performance change due to costs is

$$\Delta r^{\text{realized}} = r^{\text{cost}} - r^{\text{nocost}}.$$

This quantity captures all effects: explicit fees, any induced changes in the trade path, and the interaction between costs and portfolio constraints.

To isolate the mechanical cost component from these endogenous effects, the analysis also computes a *pure cost drag* estimate. This quantity is based on the total commissions C paid over the backtest period, distributed over the average invested capital:

$$\Delta r^{\text{pure}} \approx -\frac{C}{\bar{V}\,T},$$

where T is the sample length in years and $\bar{V}$ is proxied by the average of the starting portfolio value and the final no-cost portfolio value. By construction, $\Delta r^{\text{pure}} \leq 0$. This estimate is used in the cost decomposition figure (Figure 6.2): any discrepancy between $\Delta r^{\text{realized}}$ and Δr^{pure} reflects endogenous implementation effects rather than the direct fee burden.

The results in Section 6.4 also report the certainty equivalent (CEQ) introduced in Chapter 4. The formal definition is unchanged; Section 6.4.1 motivates reporting CEQ alongside mean returns when evaluating cost-inclusive paths.

6.4 Results

6.4.1 Strategy-level comparison

Mean return alone can be misleading when comparing no-cost and cost-inclusive runs: frictions change not only average returns but also the risk and shape of the wealth path (e.g. through cash drag, discrete sizing, and turnover). The certainty equivalent (CEQ), defined in Chapter 4, summarizes both the mean and the variance of returns for a mean–variance investor with risk aversion $\gamma = 2$. It is therefore a natural complement to Δr here: if costs depress the mean but also reduce exposure or turnover-induced volatility, CEQ indicates whether the net effect is still attractive in utility terms. The same $\gamma = 2$ and the same annualization convention as in Chapter 5 are used for comparability.

Table 6.1 reports the main outputs. Returns and CEQ are annualized (decimal form). Turnover is the volume-based annualized turnover under the cost-inclusive run (total traded notional divided by average portfolio value and the number of years); unlike the weight-change-based TO in Chapter 4, this quantity reflects actual dollar trading activity and scales directly with the commission bill. The column Δr (pp) is the realized change in annualized return in *percentage points* (pp) per year; for example, -1.70 means the cost-inclusive return is 1.70 pp/year below the no-cost return. "Fee/Final" is total commissions over the backtest as a percentage of the no-cost run's final value, $100 \times C/V_{\text{final}}^{\text{nocost}}$. "Fee/Gain" is fees relative to the gross capital gain of the no-cost run, $100 \times C/(V_{\text{final}}^{\text{nocost}} - V_0)$; see the discussion of ill-defined cases in Section 6.4.3.

Strategy	r no-cost	r cost	Δr (pp)	CEQ no-cost	CEQ cost	TO (ann.)	Fee/Final (%)	Fee/Gain (%)
EW	0.2084	0.2136	+0.52	0.0919	0.0938	0.8601	0.48	0.58
ERC	0.0792	0.0762	-0.30	0.0484	0.0464	0.6784	0.29	0.58
DRP	0.0782	0.0775	-0.08	0.0464	0.0459	0.6293	0.32	0.64
MDR	0.0566	0.0588	+0.22	0.0347	0.0361	0.4624	0.27	0.70
MDC	0.0826	0.0656	-1.70	0.0455	0.0355	1.0047	0.59	1.15
GMV	0.0355	0.0331	-0.24	0.0220	0.0204	0.3760	0.15	0.57
MSR	0.0798	0.0753	-0.46	0.0362	0.0314	2.6634	1.29	2.58
UTIL_5	0.0729	0.0573	-1.56	0.0332	0.0230	2.7415	1.23	2.62
UTIL_2	0.0826	0.1325	+5.00	0.0233	0.0387	3.2200	2.45	4.80
UTIL_0.5	-0.0009	0.1328	+13.37	-0.0755	0.0457	2.2200	4.41	*n/a*†

Table 6.1: Strategy comparison under the execution engine: no-cost vs cost-inclusive backtest. Δr is reported in percentage points (pp) per year. "Fee/Final" is the ratio of total commissions to the no-cost run's final value. "Fee/Gain" is the ratio of total commissions to the gross capital gain of the no-cost run; it becomes ill-defined when the no-cost gross gain is negative. Large positive Δr for `UTIL_2` and `UTIL_0.5` come from path-dependent exposure cuts described in Section 6.4.3, not from a claim that costs raise expected returns in general.

† `UTIL_0.5` finished the no-cost backtest below its starting value (annualized return -0.09%), so the gross capital gain is negative and the Fee/Gain ratio is undefined. Use Fee/Final (4.41%) and the CEQ columns instead.

Pattern 1: High turnover and the expected cost channel

The largest negative realized return changes are observed for MDC and the utility-based portfolios with high turnover. MDC loses about 1.7 percentage points of annual return when costs are introduced, and its CEQ falls accordingly. MSR and UTIL_5 also exhibit economically meaningful deterioration in both return and CEQ. This aligns with the broad empirical evidence that execution costs scale with trading intensity and liquidity demand (Chan and Lakonishok 1995), and with the theoretical implication that costs create effective no-trade regions that punish strategies which attempt to continuously track unstable targets (Davis and Norman 1990).

Pattern 2: Realized impact versus pure commission drag

The estimated pure commission drag is small in annualized terms for most strategies (on the order of a few tens of basis points, as shown in Figure 6.2). However, the realized Δr ranges from about -1.7 pp/year to $+13.4$ pp/year across strategies. This dispersion is too large to be explained by commissions alone and indicates that the cost-inclusive run changes the trading trajectory and, therefore, the realized return path. In other words, costs act not only as a tax on turnover but also as a *regularizer*: when the strategy is forced to internalize a trading penalty, its effective exposures may change in ways that accidentally improve or worsen performance over the sample. These positive realized cases are sample-specific outcomes and should not be interpreted as evidence that costs increase expected returns in general.

6.4.2 Visual evidence: where costs matter and why

To complement Table 6.1, Figure 6.1 plots annualized returns under the two scenarios. The figure makes it clear that costs do not merely shift all strategies down by a constant amount: for some strategies, the cost-inclusive return deviates materially from the no-cost return, which is consistent with the path dependence introduced by execution and cash constraints (Section 6.2.4).

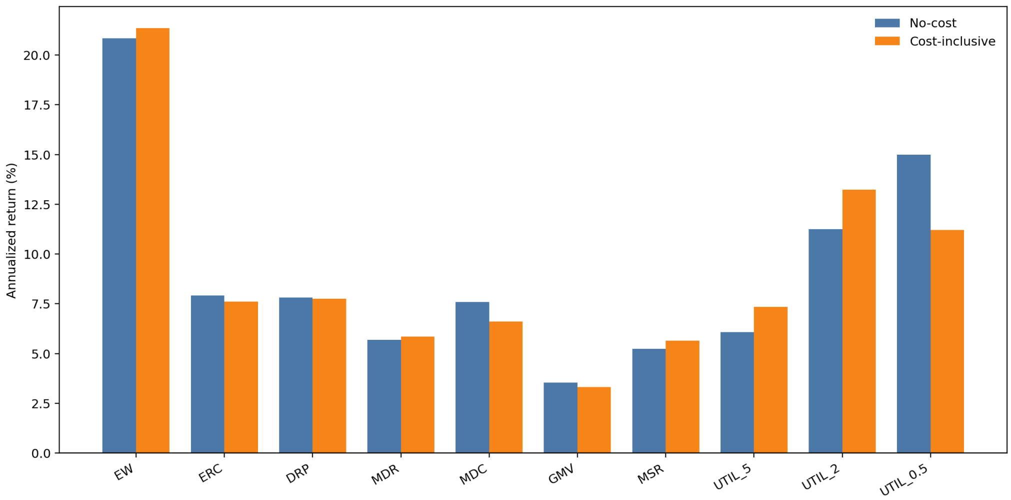

Figure 6.1: Annualized returns under the execution engine: no-cost baseline (blue) vs cost-inclusive scenario (orange). The difference between the two bars for each strategy corresponds to the realized Δr in Table 6.1.

Figure 6.2 provides a direct diagnostic of the decomposition defined in Section 6.3. For each strategy, the red bar shows the realized Δr (the full difference between cost-inclusive and no-cost returns, in percentage points per year), and the teal bar shows the pure commission drag (the annualized fee burden estimated from total commissions paid). For most strategies, the pure drag is economically small, while realized changes are sometimes substantially larger. The gap between the two bars represents the *implementation response*: the change in the realized holdings trajectory caused by execution effects.

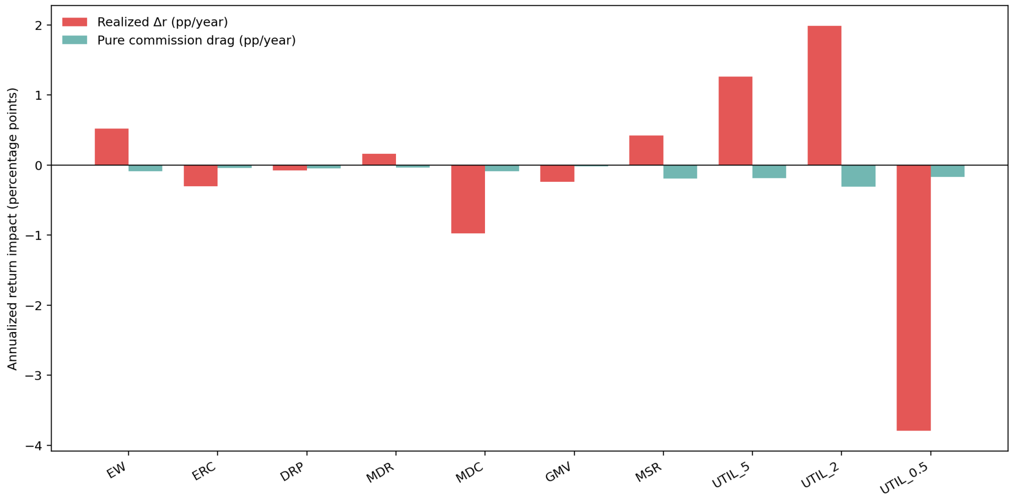

Figure 6.2: Cost impact decomposition: realized change in annualized return (red) vs pure commission drag (teal), both in percentage points per year. The gap between the two bars reflects endogenous implementation effects.

Figure 6.3 combines two views of the same comparison output: panel (a) relates annualized turnover to fee leakage (total commissions as a share of no-cost final wealth); panel (b) relates turnover to the realized change in annualized return (Δr in percentage points per year). Together they link implementability (turnover) to both the mechanical fee burden and the realized return impact. Panel (a) shows the expected positive association between trading intensity and fees; panel (b) shows that the return impact is not a simple function of turnover, consistent with path-dependent execution effects discussed in Section 6.2.4.

6.4.3 Interpreting the non-monotonic cases

Two strategies (EW and MDR) show small positive realized Δr, and the utility portfolios with $\gamma = 2$ and $\gamma = 0.5$ show large positive changes. Such outcomes should not be interpreted as evidence that trading costs “increase” expected returns. Instead, they reflect a combination of implementation effects and sampling variation:

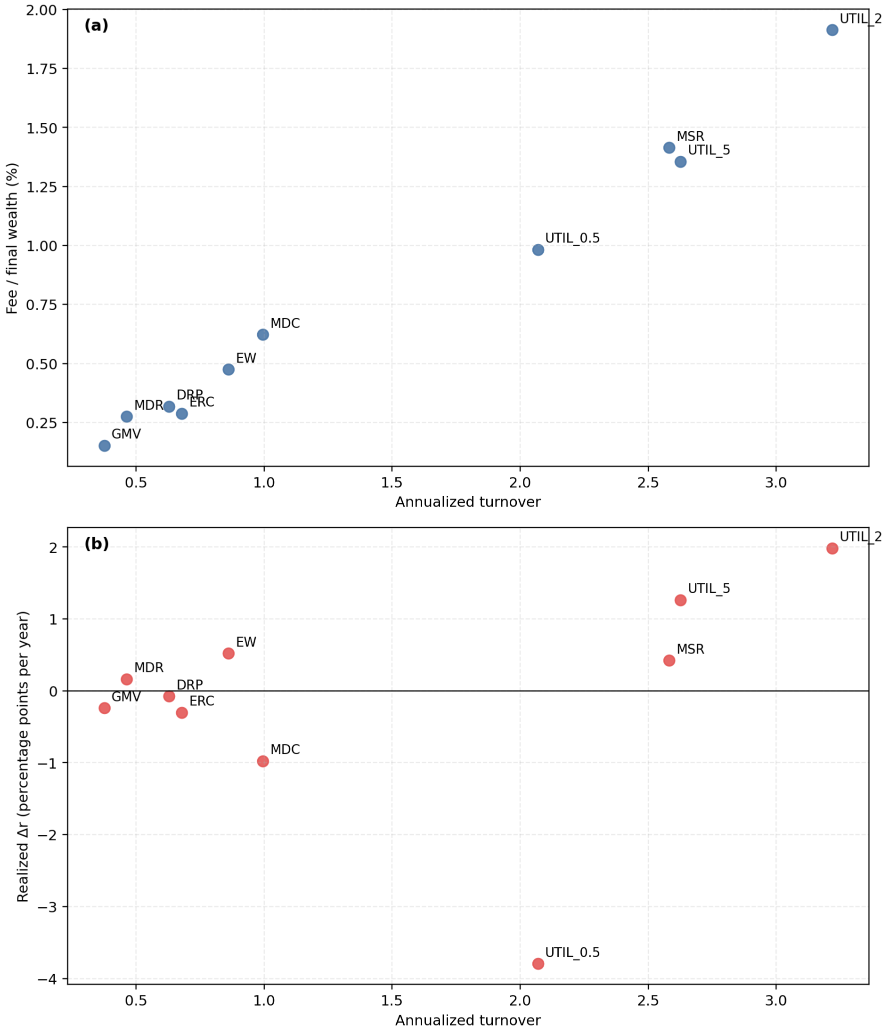

Figure 6.3: Turnover, fee leakage, and realized cost impact. **(a)** Fee paid over the full backtest as a percentage of no-cost final wealth vs annualized turnover. **(b)** Realized Δr (pp/year) vs annualized turnover. Both panels use the same strategies and the same source file as Table 6.1.

- **Endogenous de-risking**: commissions and cash drag can reduce gross exposure after rebalances, which may lower realized volatility and limit drawdowns. Over a finite sample, this can mechanically raise the annualized return computed from the wealth path, even though the strategy is paying costs.
- **Path dependence and discrete execution**: if trades occur at different prices or with slightly different sizing due to cash constraints, the sequence of holdings can diverge from the no-cost scenario. This divergence can be amplified for unstable strategies (notably the utility optimizers), which are known to be sensitive to input estimation error (Michaud 1989).
- **"Fee/Gain" when the benchmark loses money**: this ratio divides total fees by the no-cost portfolio's gross capital gain ($V_{\text{final}}^{\text{nocost}} - V_0$). It is meant to express fees as a share of profits. If the no-cost run finishes *below* its starting value, that gain is negative: the denominator has the wrong sign, so the ratio no longer answers a meaningful "fees versus profits" question. In that case, rely on **Fee/Final** (fees scaled by ending wealth) and on the change in **CEQ**.

The utility portfolios with $\gamma = 2$ and $\gamma = 0.5$ warrant closer inspection because their positive realized Δr (+5.0 and +13.4 pp/year respectively) is an order of magnitude larger than what commission drag alone could explain. The mechanism is path-dependent exposure reduction: both profiles hold aggressive crypto positions in the no-cost run, which captures the 2021 bull market but then suffers severe drawdowns in the May 2022 stress period. In the cost-inclusive run, commission drag after each rebalance slightly reduces the capital available to reinvest, which lowers effective crypto exposure incrementally over time. Over the May 2022 window this translates into a smaller realized loss than in the no-cost path, and because the 2022 drawdown is the dominant event in the sample, avoiding even a fraction of it produces a large positive Δr when annualized over nine years. This is a finite-sample coincidence, not evidence that transaction costs raise expected returns in general. The CEQ comparison (Table 6.1) is the appropriate cross-check: UTIL_2 CEQ improves from 0.0233 to 0.0387, and UTIL_0.5 from -0.0755 to 0.0457, reflecting that the cost-induced de-risking reduced variance enough to more than offset the fee burden over this specific sample.

From a thesis perspective, the key takeaway is not that costs can improve performance in expectation, but that *a purely frictionless ranking can be unstable* once a realistic execution layer is introduced. This reinforces the motivation for reporting implementability metrics (TO and TTO) alongside SR/ASR/CEQ in

Chapter 5.

6.4.4 Cost sensitivity by strategy family

The comparison in Table 6.1 becomes more interpretable once strategies are grouped by the source of their inputs and by the stability of their target weights.
Mean–variance and utility optimization. MSR and the utility portfolios depend on estimates of expected returns. As discussed in Chapter 4, expected return estimates are the noisiest inputs in portfolio optimization, and their estimation error can translate into extreme and unstable weights (Michaud 1989). The execution layer provides an economic mechanism through which this instability becomes costly: unstable targets imply larger trading volumes at each rebalance (high TTO/TO), and larger volumes imply higher fees and slippage. The result is a double fragility: statistical sensitivity to μ and economic sensitivity to trading costs.

The three utility profiles illustrate how the risk-aversion parameter γ acts as an implicit stabilizer. High risk aversion ($\gamma = 5$) tends to produce more conservative allocations and, in turn, mitigates exposure swings, yet its implementation still trades heavily in the present run. Low risk aversion ($\gamma = 0.5$) is the most aggressive profile and is therefore the most exposed to path dependence. This is precisely the regime where CEQ is informative: even when the annualized return looks favorable in the cost run, the economic interpretation requires checking whether variance and drawdowns are reduced in a meaningful way (Chapter 4, Section 4.5).
Risk-based and smart beta strategies. GMV, ERC, DRP, MDC, and MDR depend primarily on second-moment information (the covariance/correlation structure). In Chapter 5, these strategies were argued to be more stable in heavy-tailed environments because they avoid direct reliance on μ. This chapter adds a second dimension of stability: *trading robustness*. GMV and the risk-parity family (ERC/DRP) show limited sensitivity to costs, consistent with their comparatively lower turnover. By contrast, MDC exhibits a material negative realized impact in this run, which suggests that a correlation-only objective can still generate unstable targets when correlations shift rapidly (as they do in mixed universes and during stress events). This is consistent with the caution in Chapter 5 that correlation-aware optimality in a hybrid universe must be interpreted jointly with implementability.
Benchmark portfolios. EW acts as a useful baseline because it does not involve optimization. Its turnover is driven mainly by the monthly rebalancing rule rather than by parameter instability. Consequently, its fee leakage remains moderate and its performance is less sensitive to the estimation error chan-

nel highlighted above. This supports the broader empirical finance lesson that naive diversification often serves as a hard-to-beat benchmark once frictions are accounted for (DeMiguel, Garlappi, and Uppal 2009).

6.4.5 Robustness of the Chapter 5 conclusions

The central qualitative conclusion of Chapter 5, that strategy performance is regime-dependent and that robustness requires jointly considering performance and stability, remains valid. This chapter adds a practical corollary: *turnover is not merely a descriptive statistic but a transmission channel through which frictions reshape the realized outcome.*

In particular, strategies that were already flagged as turnover-heavy in the frictionless analysis (notably MSR and the return-seeking utility profiles) continue to look fragile under a realistic cost model, both in terms of net return and (more importantly) in terms of CEQ. Conversely, strategies with lower turnover, such as GMV and the risk-parity family (ERC/DRP), display much smaller sensitivity to costs, which supports their interpretation as implementation-stable baselines. This is consistent with the execution literature, where optimal trading schedules and net performance depend critically on the cost parameters and on trading intensity (Almgren and Chriss 2001).

6.4.6 Link back to regime-dependent optimality

Chapter 5's central message was that the "best" allocation rule depends on the asset universe (crypto-only, traditional-only, or mixed) and on the objective function (risk-adjusted performance versus stability and implementability). This chapter does not overturn that conclusion; rather, it *sharpens* it. Once costs are introduced, differences in turnover become economically salient, and strategies that were already fragile due to estimation error (notably the mean–variance family that depends on μ, such as MSR and low-risk-aversion utility portfolios) face an additional penalty through trading intensity.

This observation is fully consistent with the methodological discussion in Chapter 4. Michaud's critique of mean–variance optimization as an "error maximizer" (Michaud 1989) predicts unstable weights under noisy inputs, which translates into high TTO (frequent changes in targets). Under frictions, high target instability translates into real trading volume, and real trading volume translates into costs. Transaction costs therefore provide an economic mechanism linking the statistical stability metrics of Chapter 5 to out-of-sample implementability.

The most important implication for the overall thesis is that a strategy should not be evaluated on SR/ASR/CEQ alone. Even when a strategy ranks highly in a frictionless comparison, it may be *dominated after costs* if the outperformance is obtained via frequent reallocations. The practical criterion is therefore a multi-dimensional stability notion: performance must be weighed against stability (TO/TTO), particularly in the dynamic crypto universe and in the mixed regime where constituent churn and correlation shifts amplify trading requirements.

6.5 Limitations and extensions

This chapter's results should be read with two limitations in mind.

- **Cost model granularity**: an all-in basis-point model is a parsimonious proxy. Real costs are state-dependent and increase with volatility, illiquidity, and order size (Chan and Lakonishok 1995). A more granular model would estimate spreads and market impact explicitly and would distinguish between maker/taker fees and venue-specific schedules. Related work on trend-following and volatility targeting also shows that design choices in signal and risk estimation can materially reduce turnover, highlighting a direct link between model specification and implementation frictions (Baltas and Kosowski 2013).

- **Finite-sample sensitivity**: because the comparison is run over a finite historical sample, path-dependent differences can create non-monotone realized effects even when the expected effect is negative. This motivates future robustness checks such as sub-period analysis, bootstrapping, and sensitivity to rebalancing frequency and cost levels.

References

Almgren, Robert, and Neil Chriss. 2001. "Optimal Execution of Portfolio Transactions." *Journal of Risk* 3(2): 5–40.

Backtrader. 2024. *Backtrader: Python Backtesting Library for Trading Strategies*. Accessed 15 March 2024. https://www.backtrader.com/

Baltas, Nick, and Robert Kosowski. 2013. "Momentum Strategies in Futures Markets and Trend-Following Funds." SSRN Working Paper. https://doi.org/10.2139/ssrn.1968996

Binance. 2024. *Trading Fees: Spot Trading Fee Schedule*. Accessed 15 March 2024. https://www.binance.com/en/fee/schedule

Chan, Louis K. C., and Josef Lakonishok. 1995. "The Behavior of Stock Prices Around Institutional Trades." *The Journal of Finance* 50(4): 1147–1174.

https://doi.org/10.1111/j.1540-6261.1995.tb04053.x

Davis, M. H. A., and A. R. Norman. 1990. "Portfolio Selection with Transaction Costs." *Mathematics of Operations Research* 15(4): 676–713. https://doi.org/10.1287/moor.15.4.676

DeMiguel, Victor, Lorenzo Garlappi, and Raman Uppal. 2009. "Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy?" *The Review of Financial Studies* 22(5): 1915–1953.

Gârleanu, Nicolae, and Lasse Heje Pedersen. 2016. "Dynamic Portfolio Choice with Frictions." *Journal of Economic Theory* 165: 487–516. https://doi.org/10.1016/j.jet.2016.06.001

Interactive Brokers. 2024. *Stocks and ETFs Commissions: U.S. Tiered Pricing*. Accessed 15 March 2024. https://www.interactivebrokers.com/en/pricing/commissions-stocks.php

Khandoker, Mozaffar, Rafiqul Bhuyan, and Harminder Singh. 2017. "Implementation Shortfall and Transaction Costs in Equity Markets." *International Journal of Finance and Banking Research* 3(6): 91–98. https://doi.org/10.11648/j.ijfbr.20170306.12

Leland, Hayne E. 1999. *Optimal Portfolio Management with Transaction Costs and Capital Gains Taxes*. Working Paper RPF-290. Haas School of Business, University of California, Berkeley.

Michaud, Richard O. 1989. "The Markowitz Optimization Enigma: Is 'Optimized' Optimal?" *Financial Analysts Journal* 45(1): 31–42.

7 Conclusions

This thesis examined whether dominant cryptocurrencies can improve portfolio performance, diversification, and implementation stability when they are selected through economically meaningful and temporally disciplined rules. The analysis combined a dynamic cryptocurrency universe, a static traditional ETF universe, and a mixed-asset setting over 2015–2023. The central finding is that cryptocurrencies can be useful portfolio constituents, but only when the allocation rule controls estimation error, turnover, and regime dependence. Crypto exposure is not automatically diversifying. It becomes valuable when sized through risk-aware or correlation-aware objectives rather than through unstable return forecasts.

The first contribution is methodological. Instead of relying on a fixed top-N cryptocurrency list, the thesis used a rolling eligibility process based on continuous activity and relative market capitalization. This design reduces look-ahead and survivorship bias while preserving the changing structure of the crypto market. The resulting universe is investable in the sense required by portfolio construction: assets enter only after they have observable history and economically meaningful scale, and they can leave when they cease to satisfy the same criteria.

The second contribution is empirical. The crypto-only results show that return-seeking optimizers, especially the Maximum Sharpe Ratio portfolio, are fragile under heavy tails and regime shifts. This is consistent with the estimation-error critique of mean-variance optimization: noisy expected returns are amplified into unstable weights. Risk-based rules, particularly GMV in the crypto-only universe, produce more defensible outcomes because they depend less on sample mean estimates. In the traditional-only universe, the ranking changes: value weighting benefits from stable equity exposure and lower input noise. The mixed universe provides the most informative result. Maximum Decorrelation emerges as the strongest compromise because it uses cross-sleeve correlation structure without letting the crypto sleeve dominate total risk. Petukhina et al. (2021) supply the closest published benchmark for multi-rule crypto portfolios. Their out-of-sample crypto-only patterns match the ranking here: when mean estimates are noisy, return-seeking optimizers deteriorate sharply and risk-based rules stay more orderly. This thesis stresses three extensions they treat less centrally: a longer 2015–2023 window, a rotating eligibility filter written to limit look-ahead and survivorship bias, and the costed backtest in Chapter 6 on the same strategy menu. On mixed traditional-plus-crypto sleeves the conclusions diverge: where they still leave room for mean–variance designs to compete, the evidence here favors Maximum Decorrelation once heavy tails and turnover-driven costs enter the same exercise.

The third contribution is operational. Chapter 6 shows that transaction costs, slippage, cash accounting, and discrete execution can alter realized returns in ways that are not captured by frictionless portfolio arithmetic. Turnover is a useful proxy for implementation risk, but the realized effect of costs is also path dependent. Strategies with unstable targets are more exposed to this implementation layer, while more stable risk-based rules preserve the main Chapter 5 conclusions more effectively.

The overall message is regime-dependent. No allocation rule dominates all universes and all metrics. GMV is well suited to the dynamic crypto-only setting, value weighting is competitive in the static traditional benchmark, and Maximum Decorrelation is the most stable rule when crypto and traditional assets are optimized jointly. For investors, the practical implication is that dominant cryptocurrencies should not be added mechanically. They should be introduced through explicit selection filters, conservative risk controls, and implementation-aware rebalancing rules.

Several limitations remain. The empirical design uses historical daily data and monthly rebalancing, so results may differ under intraday execution, alternative liquidity filters, or different cost assumptions. The study also excludes stablecoins, wrapped assets, and meme-driven categories to focus on risky capital-appreciation assets, but these segments may matter for other research questions such as cash management or decentralized finance collateral. Future work could extend the analysis to adaptive thresholds, regime-switching covariance models, explicit liquidity constraints, and out-of-sample tests after the introduction of spot bitcoin exchange-traded products in 2024. These extensions would further clarify when digital assets provide genuine diversification rather than simply adding another source of concentrated risk.

The code and configuration files used to reproduce the data pipeline, portfolio simulations, and figures are available at `https://github.com/landifrancesco/Master-Thesis`.

A Diagnostic plots

This appendix collects, for each asset i, a single composite figure with four diagnostics:

1. daily log-returns r_t,
2. 30-day rolling volatility $\hat{\sigma}_{t,30}$ (std of r_t),
3. sample ACF of returns $\rho_k(r_t)$,
4. sample ACF of squared returns $\rho_k(r_t^2)$.

For items 3 and 4, lag 0 is omitted in the plots.

ADA – Cardano

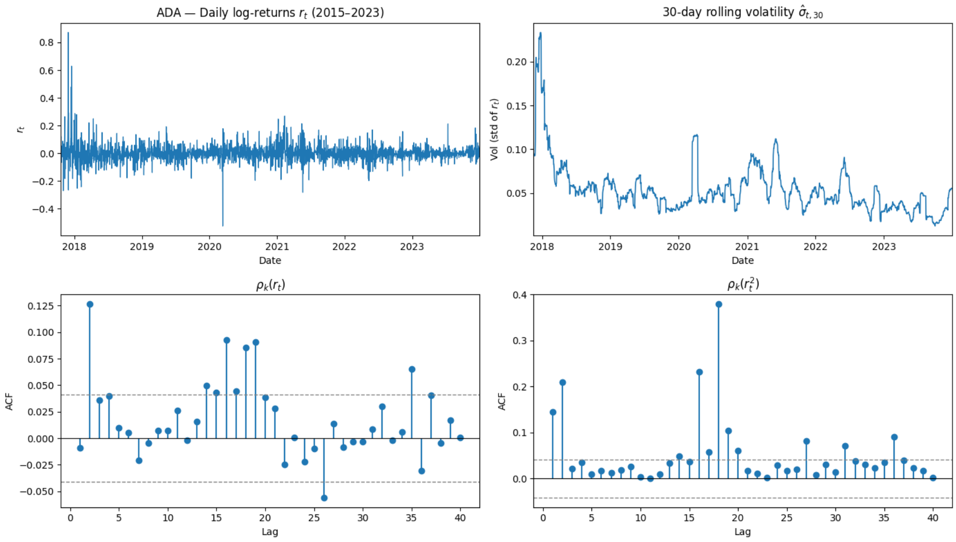

Appendix Figure A.1: Diagnostic plots for `ADA` (Cardano).

- Heavy-tailed r_t with intermittent large shocks.
- Persistent clustering in $\hat{\sigma}_{t,30}$.
- $\rho_k(r_t) \approx 0$ for $k \geq 2$.
- $\rho_k(r_t^2) > 0$ at small k with slow decay.

AVAX – Avalanche

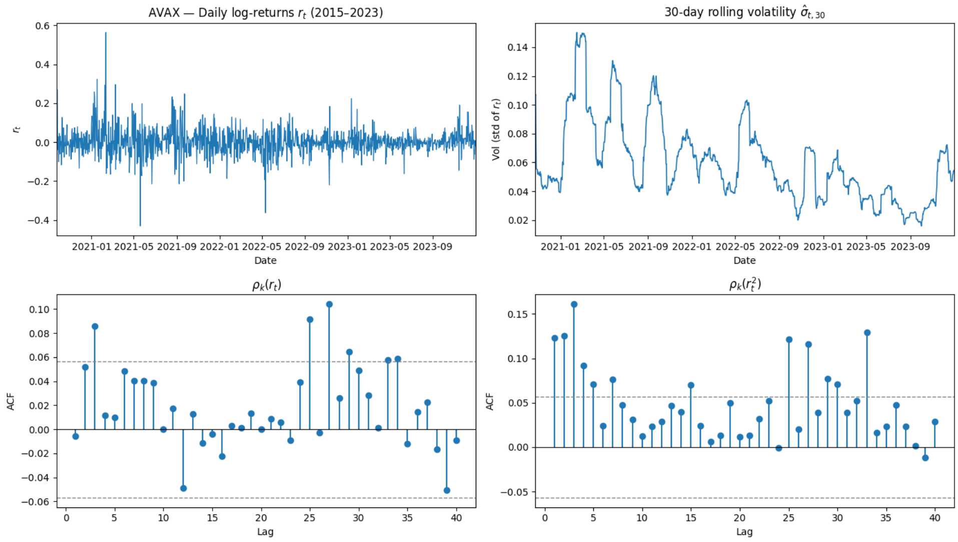

Appendix Figure A.2: Diagnostic plots for `AVAX` (Avalanche).

- Pronounced regime switches in $\hat{\sigma}_{t,30}$.
- Large bilateral excursions in r_t.
- $\rho_k(r_t) \to 0$ rapidly.
- $\rho_k(r_t^2)$ positive and slowly decaying at low k.

BCH – Bitcoin Cash

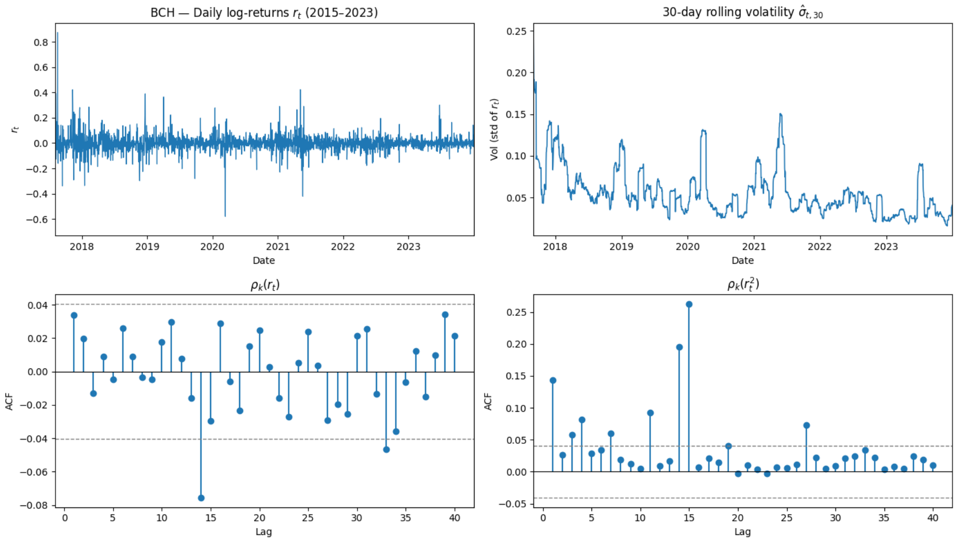

Appendix Figure A.3: Diagnostic plots for `BCH` (Bitcoin Cash).

- Event-driven spikes in r_t.
- Medium–high volatility clustering.
- $\rho_k(r_t) \approx 0$ beyond the first lag.
- Short-memory but strong dependence in r_t^2.

BNB – BNB

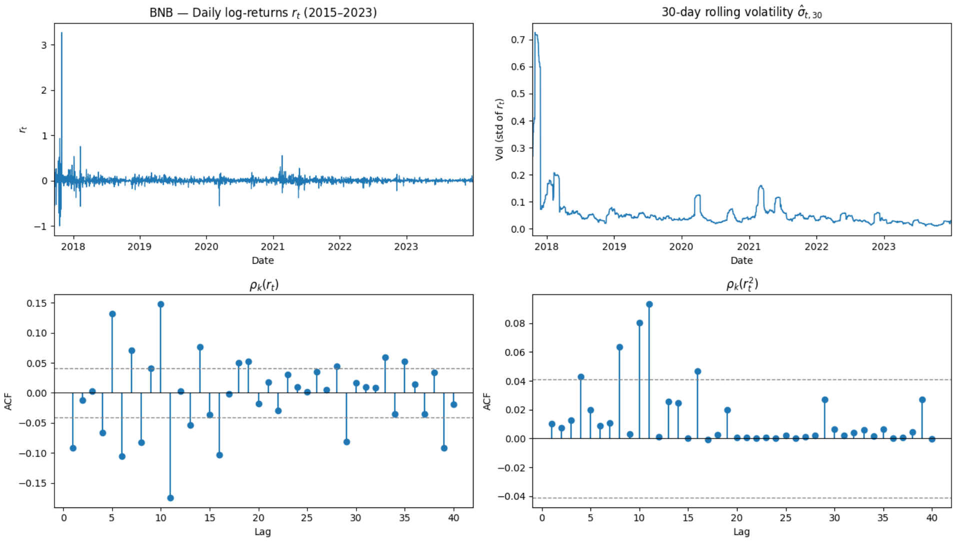

Appendix Figure A.4: Diagnostic plots for `BNB`.

- Extreme tail events; marked leptokurtosis in r_t.
- Very high and persistent $\hat{\sigma}_{t,30}$ during episodes.
- $\rho_k(r_t) \approx 0$.
- Large, slowly decaying $\rho_k(r_t^2)$.

BSV – Bitcoin SV

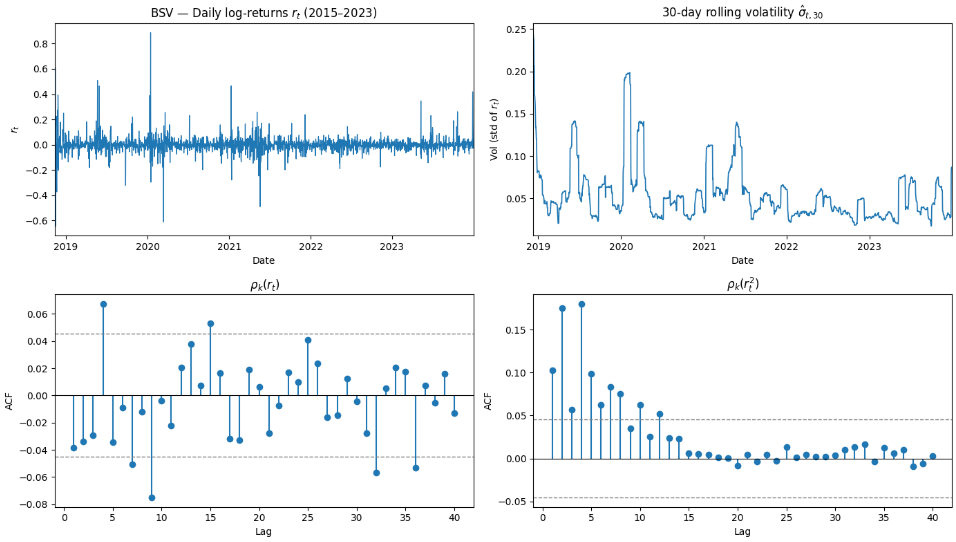

Appendix Figure A.5: Diagnostic plots for `BSV` (Bitcoin SV).

- Episodic jumps in r_t.
- Block-wise clustering in $\hat{\sigma}_{t,30}$.
- $\rho_k(r_t) \approx 0$ beyond minimal lags.
- Strong short-lag $\rho_k(r_t^2)$.

BTC – Bitcoin

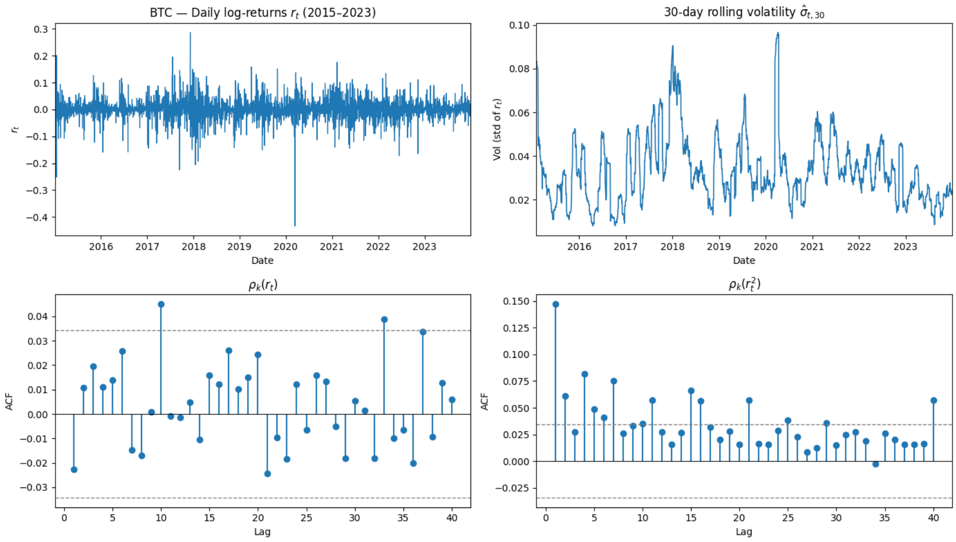

Appendix Figure A.6: Diagnostic plots for BTC (Bitcoin).

- Lower dispersion relative to altcoins.
- Clear but milder clustering in $\hat{\sigma}_{t,30}$.
- $\rho_k(r_t) \approx 0$.
- $\rho_k(r_t^2)$ significant at small k with faster decay.

CRO – Cronos

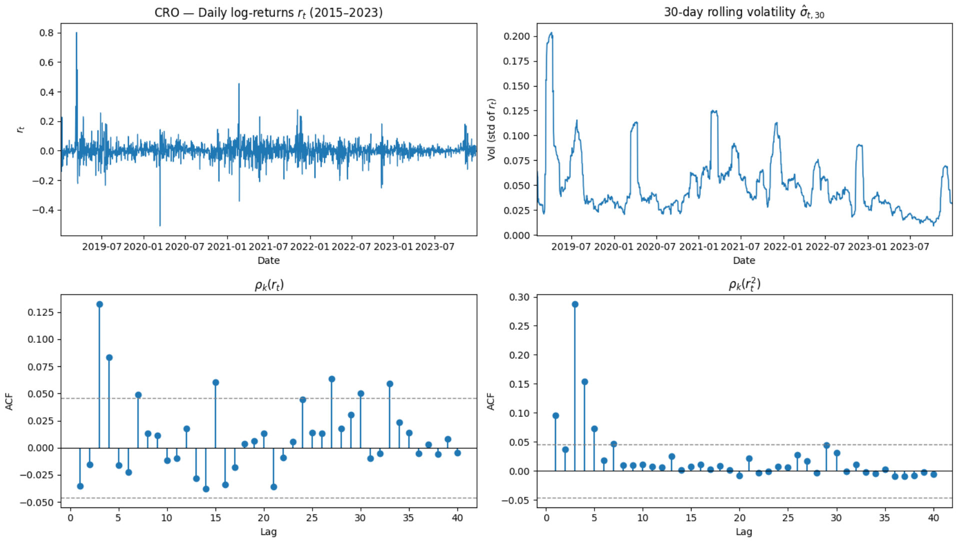

Appendix Figure A.7: Diagnostic plots for `CRO` (Cronos).

- Volatility compression–expansion cycles.
- Heavy-tailed r_t.
- $\rho_k(r_t) \approx 0$.
- Persistent $\rho_k(r_t^2)$ at low k.

DOT – Polkadot

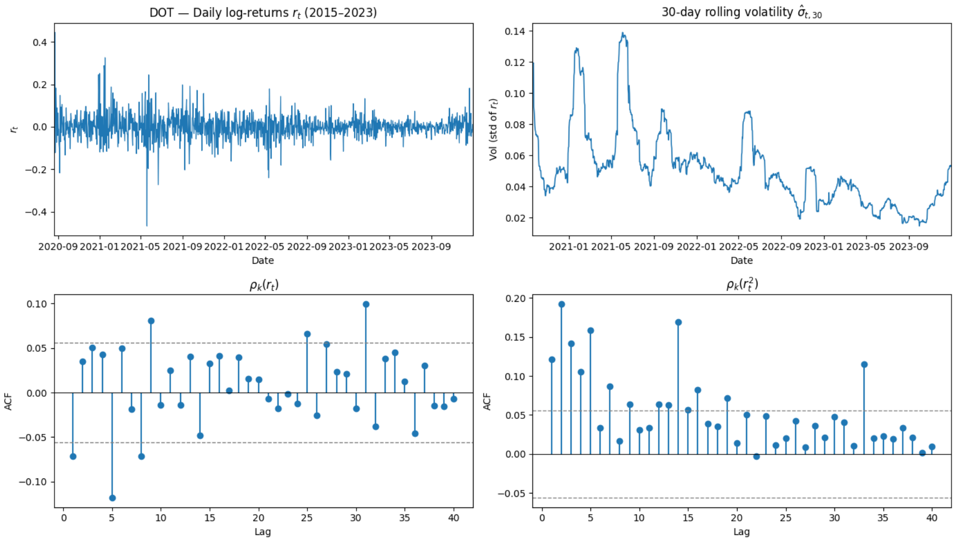

Appendix Figure A.8: Diagnostic plots for `DOT` (Polkadot).

- Well-separated volatility regimes.
- Dispersed extremes in r_t.
- $\rho_k(r_t) \to 0$ quickly.
- Slowly decaying $\rho_k(r_t^2)$.

EOS – EOS

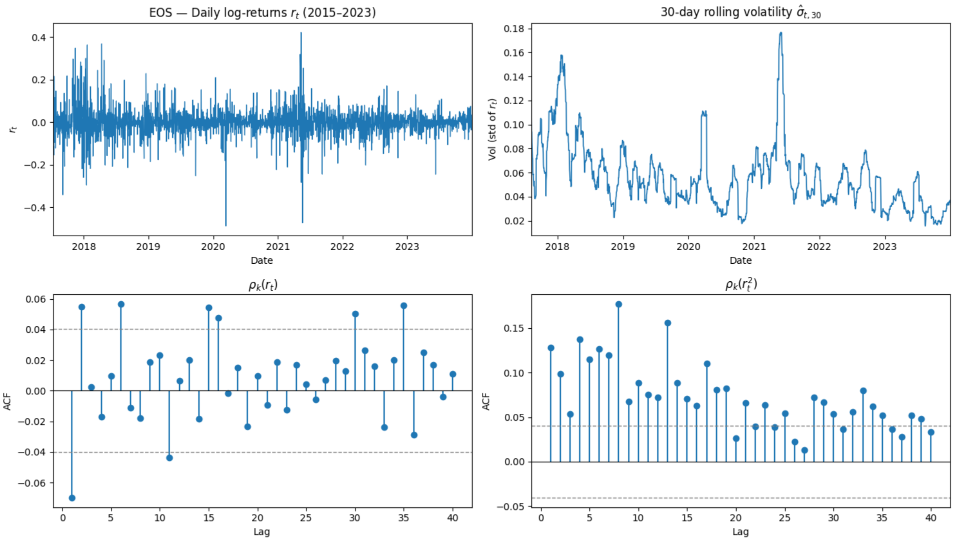

Appendix Figure A.9: Diagnostic plots for EOS.

- Concentrated shock windows followed by normalization of $\hat{\sigma}_{t,30}$.
- Leptokurtic r_t.
- $\rho_k(r_t) \approx 0$, $\rho_k(r_t^2) > 0$ for small k.

ETC – Ethereum Classic

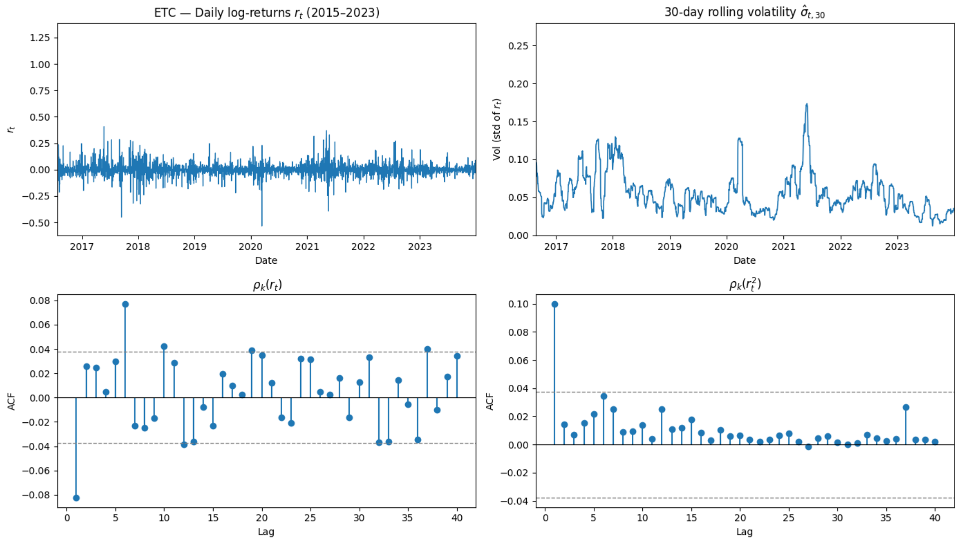

Appendix Figure A.10: Diagnostic plots for `ETC` (Ethereum Classic).

- Rare, very large outliers in r_t.
- Long volatility spells.
- $\rho_k(r_t) \approx 0$.
- Elevated and persistent $\rho_k(r_t^2)$.

ETH – Ethereum

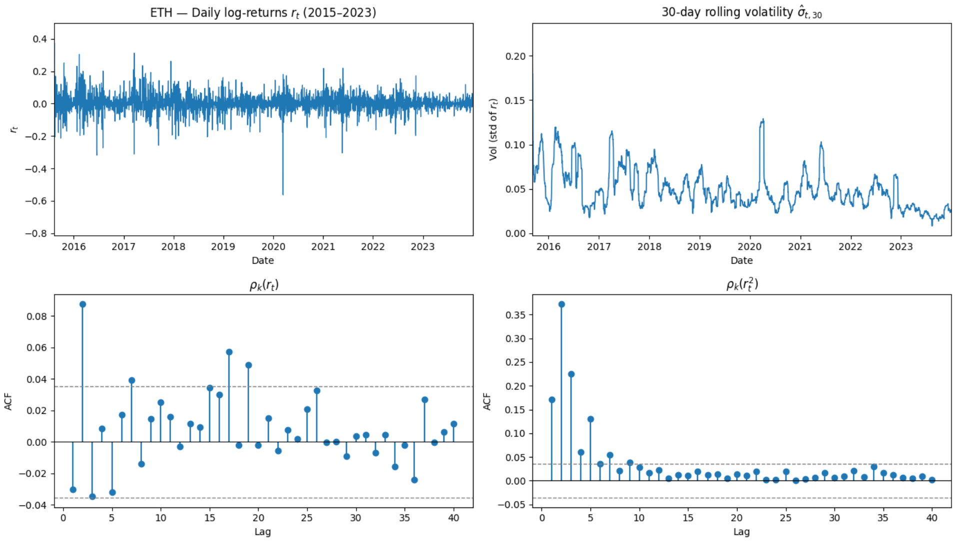

Appendix Figure A.11: Diagnostic plots for `ETH` (Ethereum).

- Higher dispersion than BTC.
- Pronounced clustering in $\hat{\sigma}_{t,30}$.
- Non-Gaussian r_t; $\rho_k(r_t) \approx 0$.
- $\rho_k(r_t^2)$ significant at short lags.

IOTA – IOTA

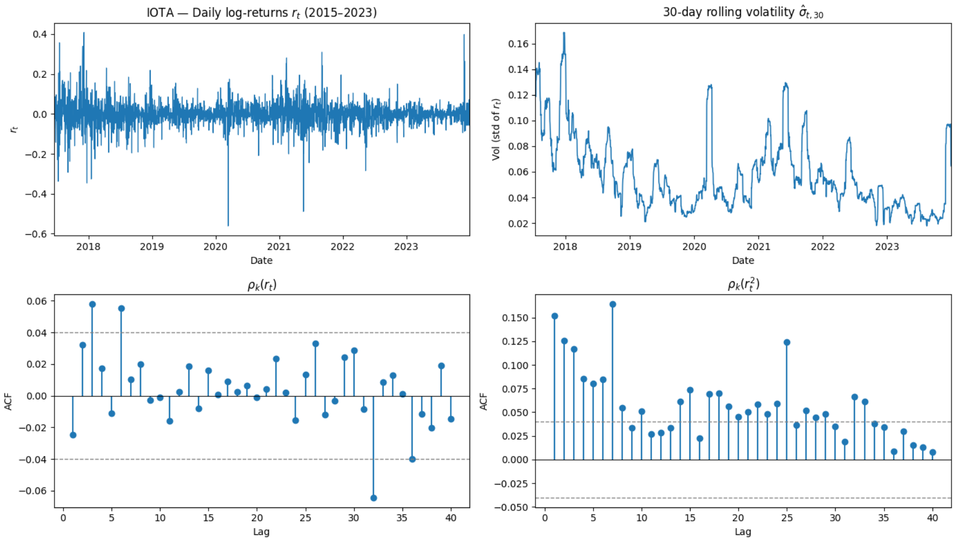

Appendix Figure A.12: Diagnostic plots for `IOTA`.

- Alternation of calm and explosive phases.
- Tail risk evident in r_t.
- $\rho_k(r_t) \approx 0$.
- $\rho_k(r_t^2) > 0$ at small k.

LINK – Chainlink

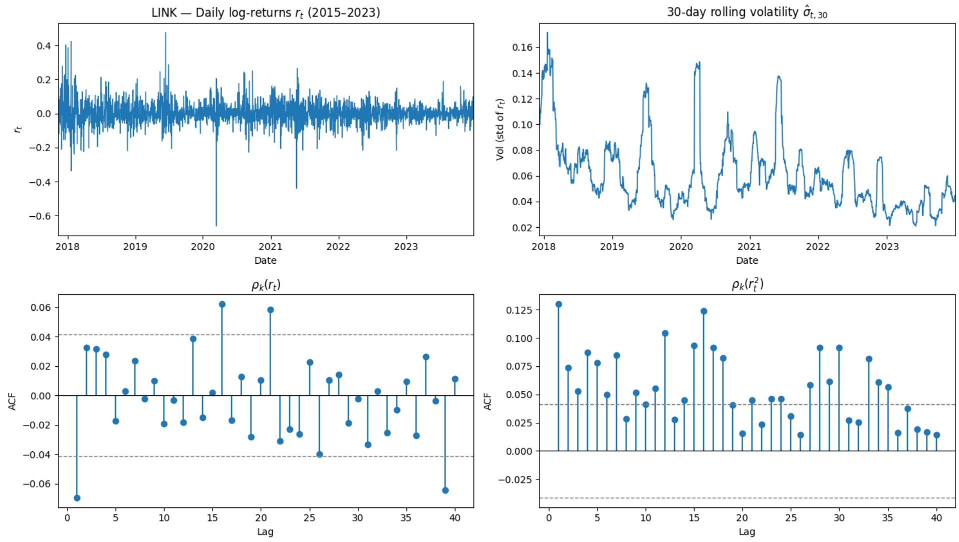

Appendix Figure A.13: Diagnostic plots for `LINK` (Chainlink).

- Period-specific surges in r_t.
- Moderate tail asymmetry.
- Volatility clustering evident.
- $\rho_k(r_t)$ negligible; $\rho_k(r_t^2)$ decays slowly.

LTC – Litecoin

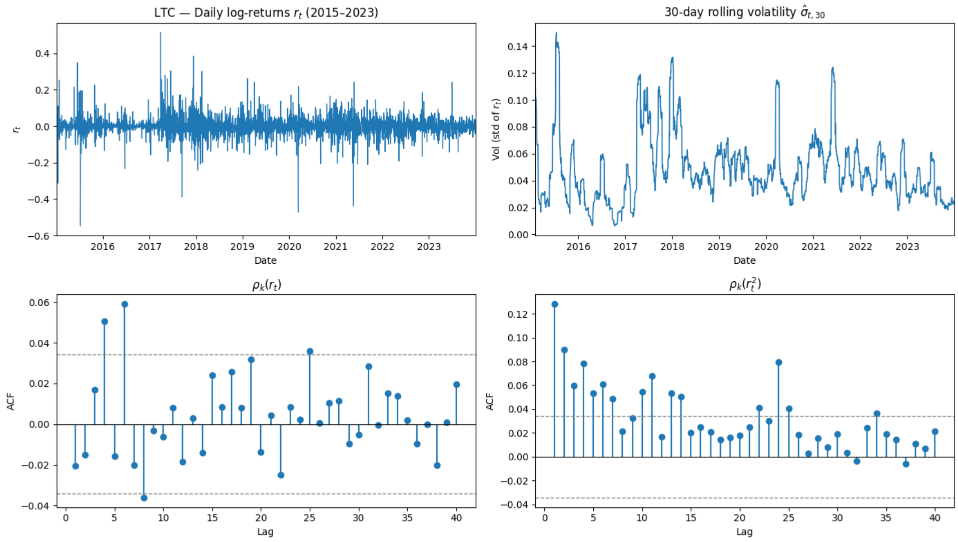

Appendix Figure A.14: Diagnostic plots for `LTC` (Litecoin).

- Cyclical volatility with fewer extremes than newer alts.
- $\rho_k(r_t) \approx 0$ beyond minimal lags.
- $\rho_k(r_t^2)$ significant at small k.

LUNC – Terra Luna Classic

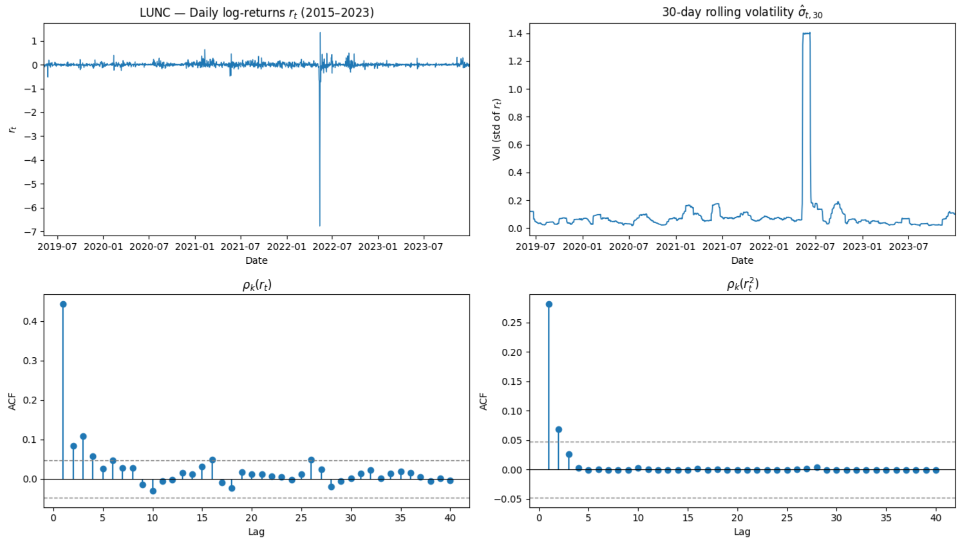

Appendix Figure A.15: Diagnostic plots for `LUNC` (Terra Luna Classic).

- Catastrophic tail event dominates r_t.
- $\hat{\sigma}_{t,30}$ spikes to extreme levels.
- $\rho_k(r_t)$ near 0 beyond the first lag.
- Very strong $\rho_k(r_t^2)$ at small k.

MATIC – Polygon

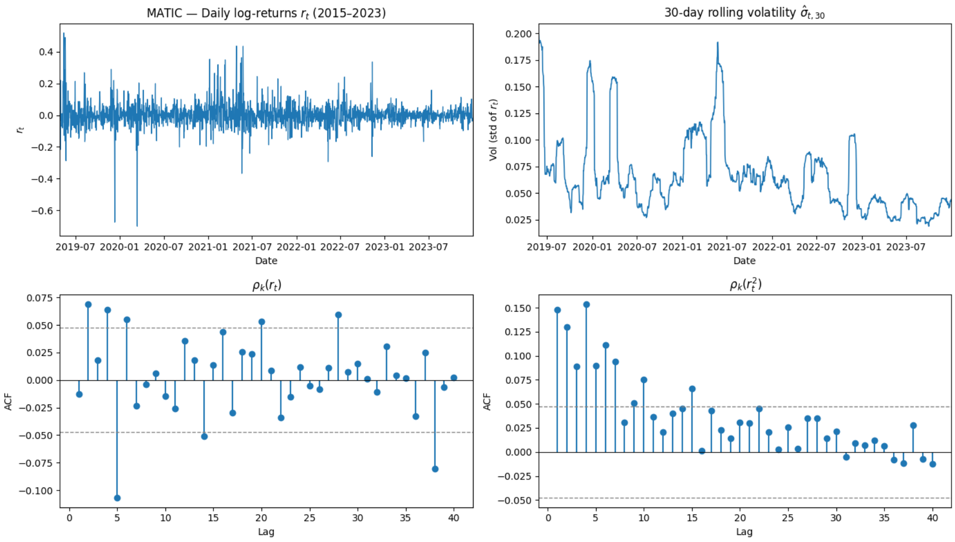

Appendix Figure A.16: Diagnostic plots for `MATIC` (Polygon).

- Sharp volatility regime transitions.
- Heavy-tailed r_t.
- $\rho_k(r_t) \approx 0$.
- Persistent $\rho_k(r_t^2)$ at low k.

NEO – NEO

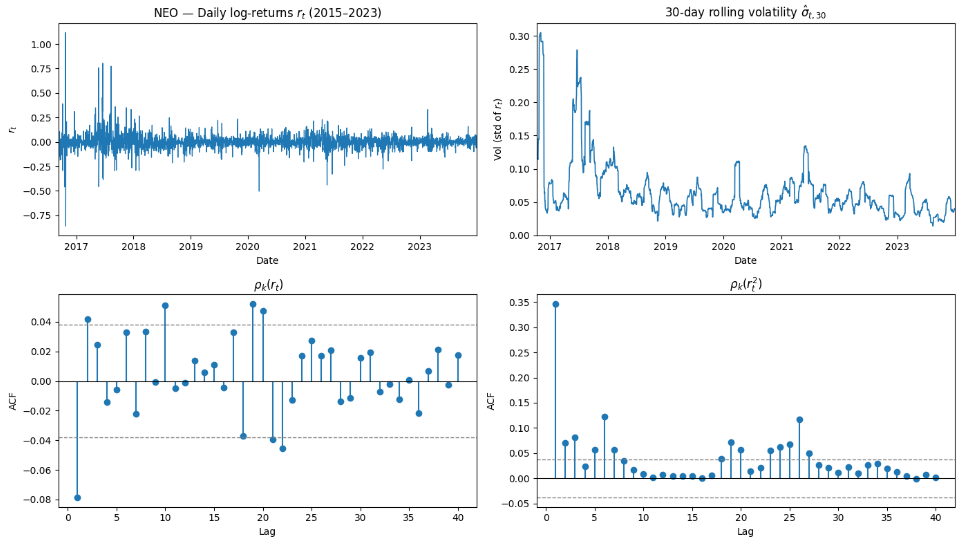

Appendix Figure A.17: Diagnostic plots for `NEO`.

- Few concentrated high-volatility periods.
- Prolonged normalization phases in $\hat{\sigma}_{t,30}$.
- $\rho_k(r_t) \approx 0$.
- Short-lag $\rho_k(r_t^2)$ significant.

SOL – Solana

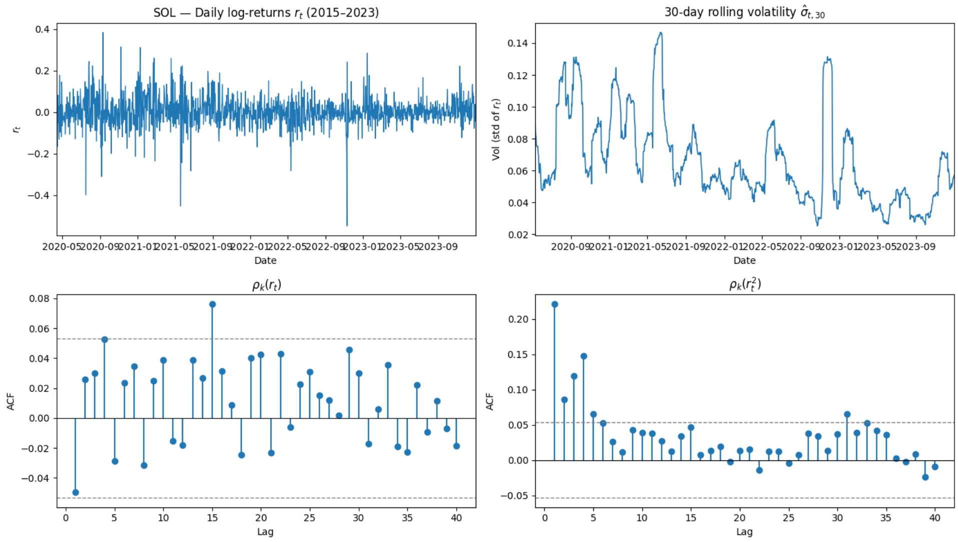

Appendix Figure A.18: Diagnostic plots for `SOL` (Solana).

- Wide regime shifts and large bilateral jumps in r_t.
- Strong clustering in $\hat{\sigma}_{t,30}$.
- $\rho_k(r_t) \approx 0$.
- Slowly decaying $\rho_k(r_t^2)$.

TRX – Tron

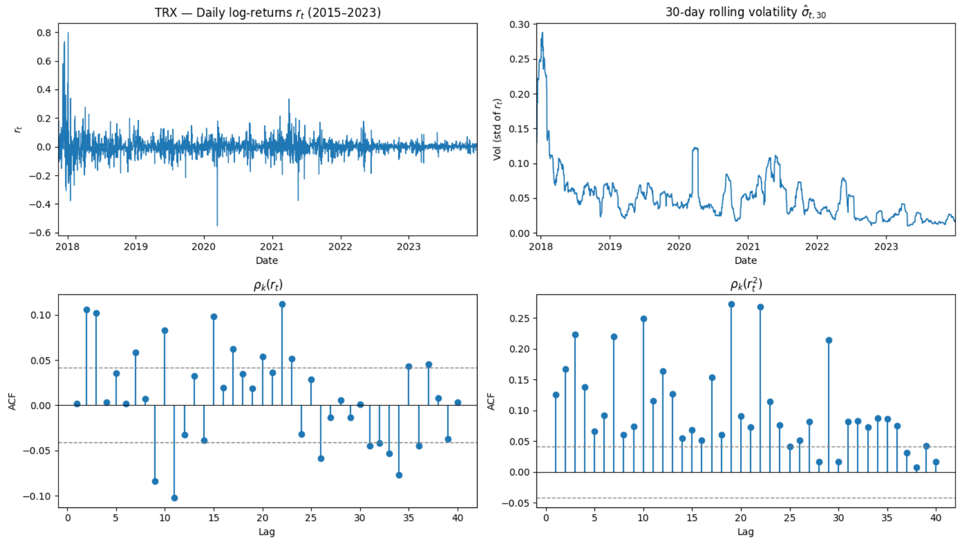

Appendix Figure A.19: Diagnostic plots for `TRX` (Tron).

- Frequent moderate outliers in r_t.
- Repeated clustering episodes.
- Negligible $\rho_k(r_t)$.
- Significant $\rho_k(r_t^2)$ at low k.

XLM – Stellar

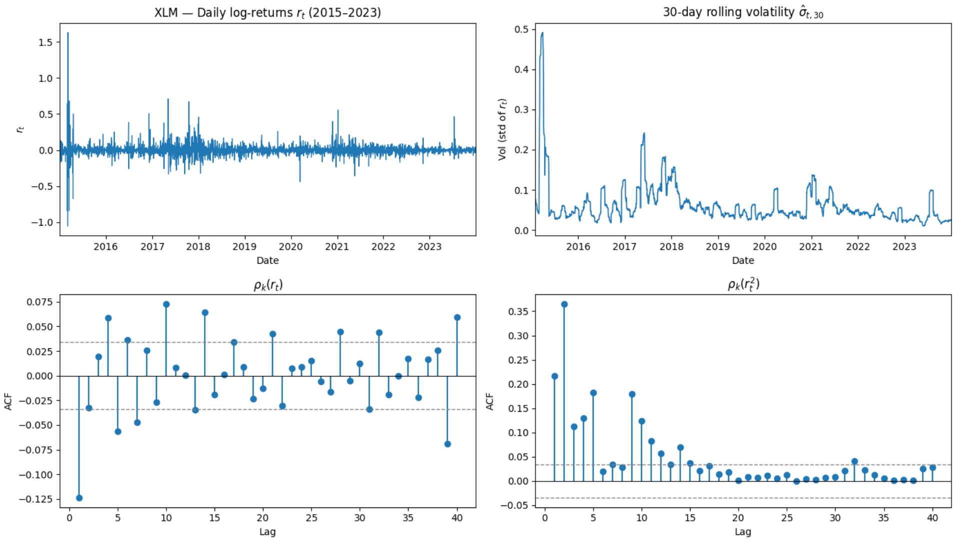

Appendix Figure A.20: Diagnostic plots for `XLM` (Stellar).

- High tail risk with very large outliers.
- Long-lived clustering in $\hat{\sigma}_{t,30}$.
- $\rho_k(r_t) \approx 0$.
- Large short-lag $\rho_k(r_t^2)$.

XRP – XRP

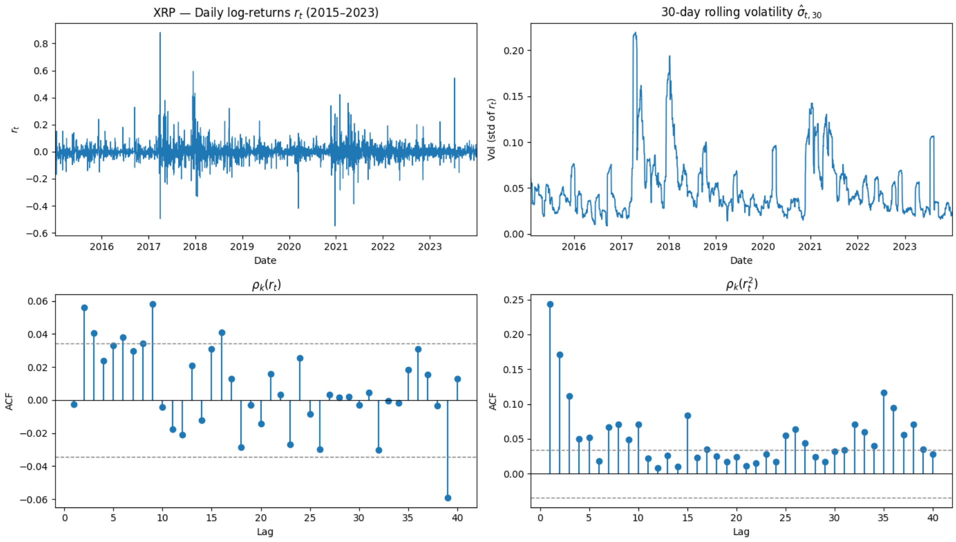

Appendix Figure A.21: Diagnostic plots for XRP.

- Event-driven bursts in r_t.
- Short memory in levels, strong short memory in volatility.
- $\rho_k(r_t) \approx 0$.
- $\rho_k(r_t^2)$ significant at low k.

DBA – Invesco DB Agriculture Fund

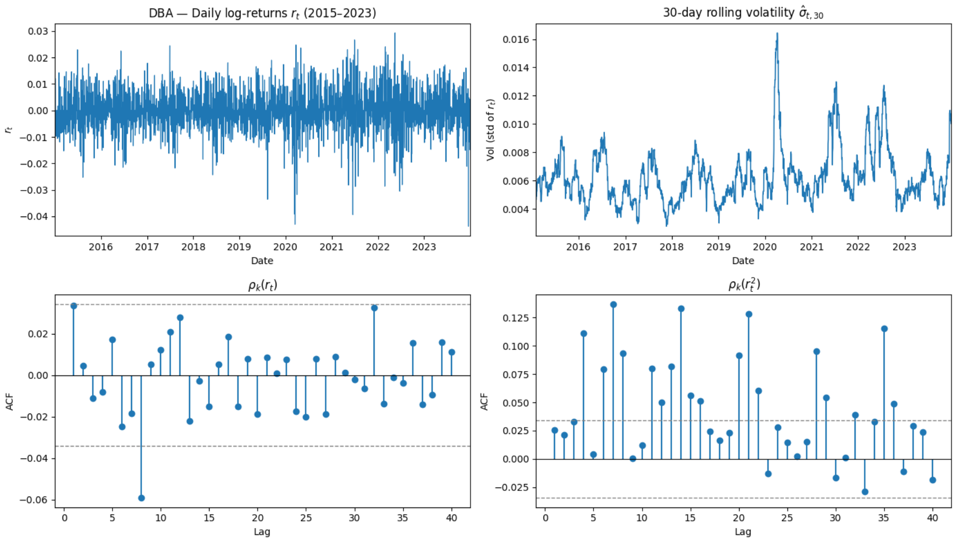

Appendix Figure A.22: Diagnostic plots for `DBA` (Invesco DB Agriculture Fund).

- Moderate dispersion of r_t with occasional commodity-driven shocks.
- $\hat{\sigma}_{t,30}$ shows clear but not extreme clustering, aligned with crop seasonality and risk episodes.
- $\rho_k(r_t) \approx 0$ beyond minimal lags.
- $\rho_k(r_t^2)$ positive at short lags, decaying moderately.

DBB – Invesco DB Base Metals Fund

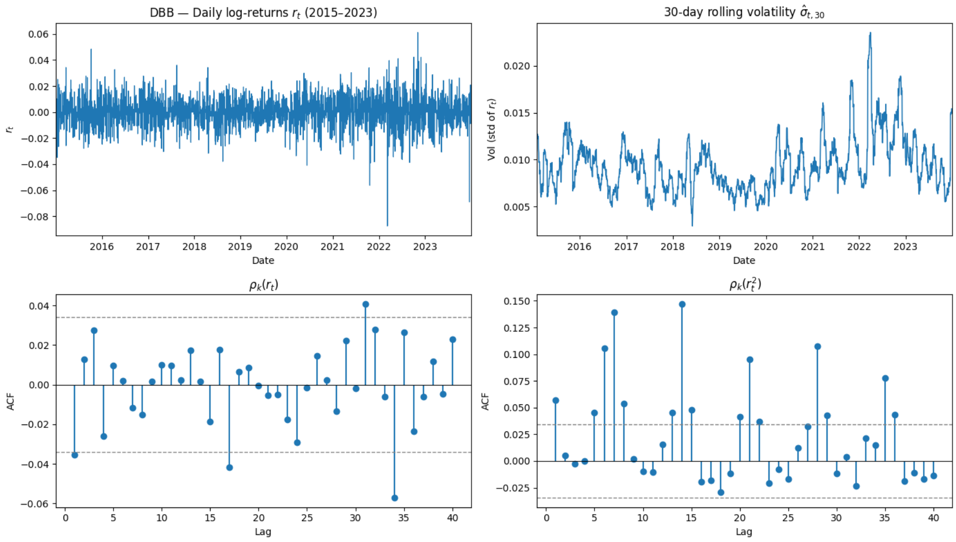

Appendix Figure A.23: Diagnostic plots for `DBB` (Invesco DB Base Metals Fund).

- High dispersion in r_t relative to agriculture and bonds.
- Noticeable volatility regimes in $\hat{\sigma}_{t,30}$ during commodity cycles.
- $\rho_k(r_t)$ close to zero after the first lag.
- $\rho_k(r_t^2)$ clearly positive at small k, with a slower decay than DBA.

DBE – Invesco DB Energy Fund

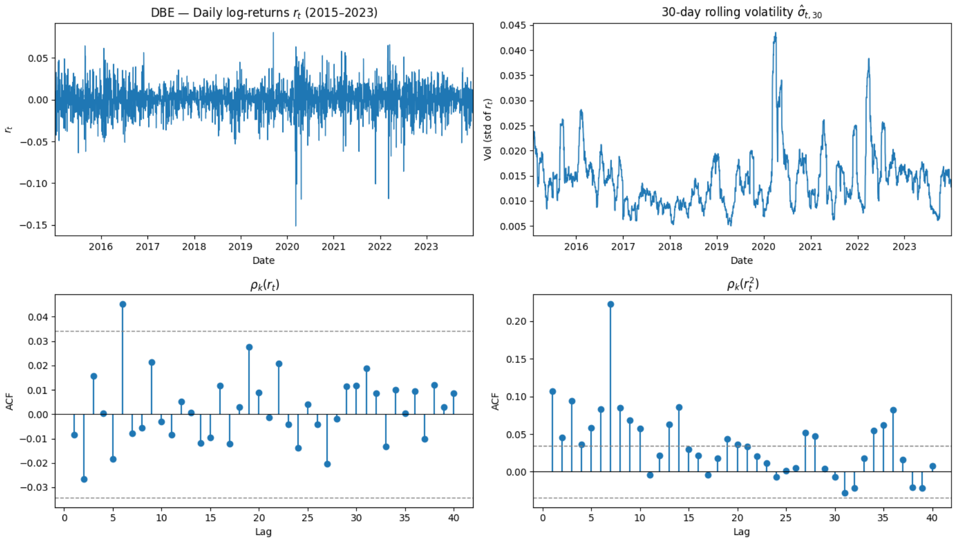

Appendix Figure A.24: Diagnostic plots for `DBE` (Invesco DB Energy Fund).

- Largest dispersion among the set; heavy downside excursions in r_t.
- Strong and persistent clustering in $\hat{\sigma}_{t,30}$ across energy stress episodes.
- $\rho_k(r_t)$ negligible beyond the first lag.
- $\rho_k(r_t^2)$ large at short lags with slow decay, indicating pronounced volatility clustering.

GLD – SPDR Gold Shares

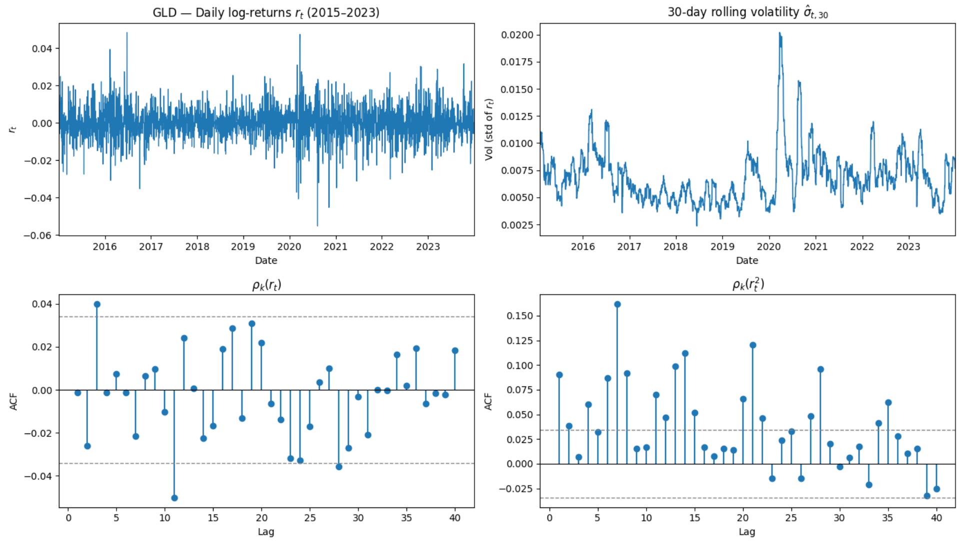

Appendix Figure A.25: Diagnostic plots for `GLD` (SPDR Gold Shares).

- Medium dispersion of r_t, roughly symmetric around zero.
- $\hat{\sigma}_{t,30}$ shows distinct but milder clustering than broad commodities.
- $\rho_k(r_t) \approx 0$ beyond minimal lags.
- $\rho_k(r_t^2)$ positive at small k, with moderate persistence.

GOVT – iShares U.S. Treasury Bond ETF

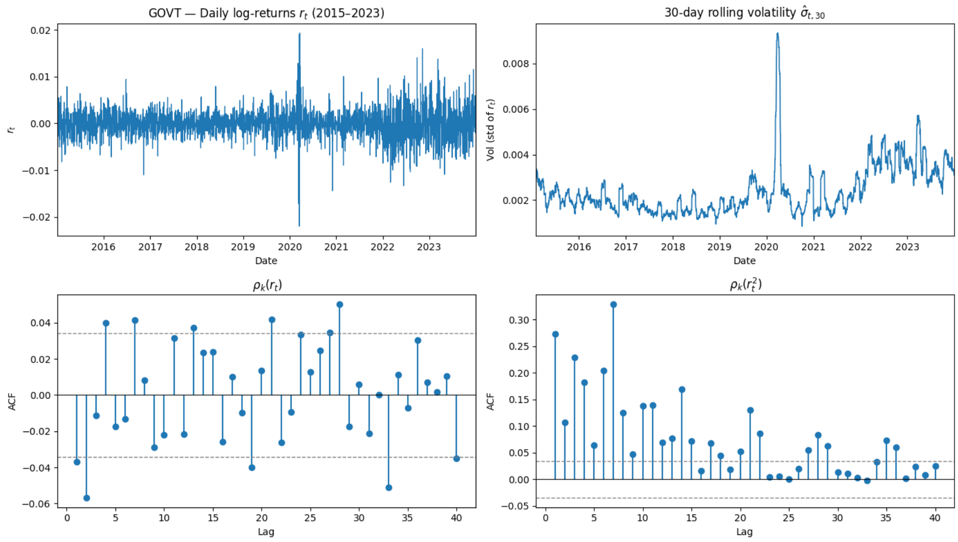

Appendix Figure A.26: Diagnostic plots for `GOVT` (iShares U.S. Treasury Bond ETF).

- Lowest dispersion of r_t; movements concentrated near zero with occasional jumps.
- $\hat{\sigma}_{t,30}$ mostly stable, punctuated by brief spikes in stress periods.
- $\rho_k(r_t)$ indistinguishable from zero across lags.
- $\rho_k(r_t^2)$ mildly positive at short lags, decaying quickly.

SPY – SPDR S&P 500 ETF Trust

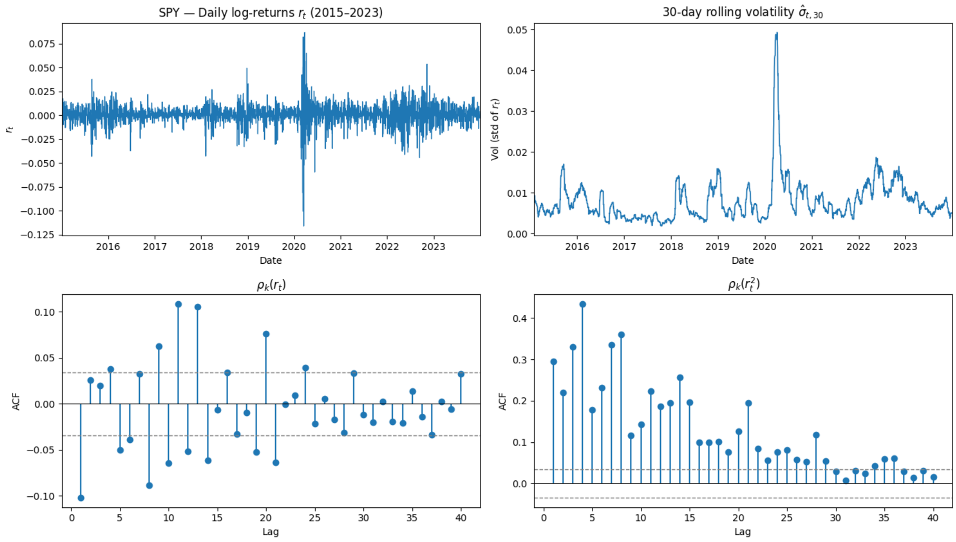

Appendix Figure A.27: Diagnostic plots for SPY (SPDR S&P 500 ETF Trust).

- Medium–high dispersion in r_t relative to bonds; rare but large downside shocks.
- $\hat{\sigma}_{t,30}$ exhibits regime shifts around market stress events.
- $\rho_k(r_t) \approx 0$ beyond the first lag.
- $\rho_k(r_t^2)$ clearly positive at small k, with persistence consistent with equity volatility clustering.

USIG – iShares Broad USD Investment Grade Corporate Bond ETF

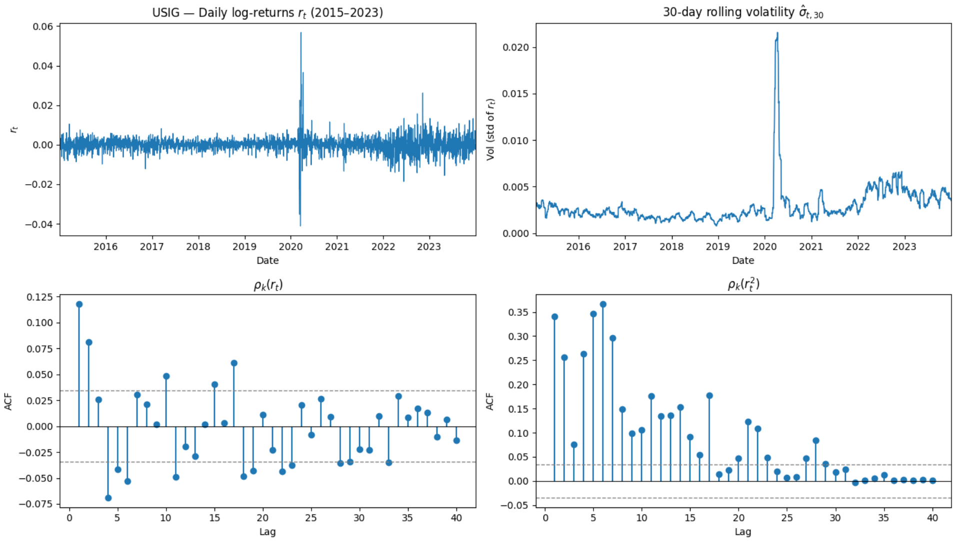

Appendix Figure A.28: Diagnostic plots for `USIG` (iShares Broad USD Investment Grade Corporate Bond ETF).

- Low dispersion in r_t overall, with sporadic large moves during credit stress and relief rallies.
- $\hat{\sigma}_{t,30}$ mostly subdued, rising noticeably in stress windows.
- $\rho_k(r_t)$ near zero for $k \geq 1$.
- $\rho_k(r_t^2)$ positive at short lags but less persistent than equities/commodities.

B Study case: the May 2022 Terra/LUNA collapse

This appendix is optional reading: it documents a single stress episode referenced in Chapters 3 and 5 and can be skipped without losing the cross-scenario argument in Chapter 5.

The most extreme single event in the 2015–2023 sample is the May 2022 collapse of the Terra ecosystem, which produced both a near-total wipeout of Terra Classic (LUNC) and the de-peg of its algorithmic stablecoin TerraUSD (UST). Because LUNC sits inside the same *Infrastructure & Interoperability* macro-category that defines the dynamic crypto sleeve in Section 5.2, the event is not an exogenous backdrop: it enters the eligible universe directly through the rolling 1% market-cap filter and, absent intervention, would have been carried by the optimisers into the portfolio. What follows isolates the episode so that its disproportionate influence on the empirical results can be read separately from the regime-level findings.

Event timeline and economic mechanism. Terra was a layer-1 blockchain whose native token LUNA was used to absorb volatility for UST, an algorithmic stablecoin pegged to one US dollar via a mint-and-burn mechanism between the two assets. Between 7–12 May 2022, a combination of large UST withdrawals from the Anchor lending protocol and concentrated sell pressure in low-liquidity venues broke the peg; the arbitrage mechanism then minted LUNA in increasing quantities to defend it, triggering a hyperinflationary spiral in which LUNA's circulating supply expanded by several orders of magnitude in a few days. By 13 May 2022 UST was trading well below $0.10 and LUNA had lost more than 99% of its value, with on-chain forensic analyses tracing the de-peg dynamics in detail (Briola et al. 2022; Nansen 2022). A subsequent fork separated the post-collapse chain (Terra 2.0, LUNA) from the original asset, which was renamed Terra Classic (LUNC) and is the version that survives in the historical CoinGecko price series used in this thesis. The U.S. Securities and Exchange Commission later filed civil charges against Terraform Labs and its founder Do Kwon, alleging a multi-billion-dollar fraud built around both UST and the broader Terra ecosystem (U.S. Securities and Exchange Commission 2023).

Impact on the dynamic crypto universe. Until April 2022 LUNA satisfied the survivorship and 1%-of-category market-cap filters defined in Section 3.5.1 and was therefore part of the eligible set used by every optimiser in Scenario A and Scenario C. Its weight in return-seeking allocations (MSR and the low-γ utility portfolios) was non-negligible in the months immediately preceding the collapse, because the asset combined a sharply rising mean return with a volatility profile that, ex post, severely understated tail risk. When the collapse materialised, daily returns reached the -99.9% magnitude visible in the descriptive statistics of Table 3.7 (the LUNC row reports a minimum daily return of -0.9989 and a kurtosis above 200), which is the exact pattern that Michaud (1989) identifies as the extreme realisation of the error-maximisation problem in mean-variance optimisation. The event therefore acts as a large, identifiable contributor to the catastrophic out-of-sample behaviour of MSR and Return-Seeking allocations documented in section 5.2, rather than being a generic source of background volatility. Figure B.1 reports an overview of the affected portfolios across strategies during the crash window.

Treatment in the backtest and robustness check. For the reasons set out in Section 3.5.1 and Chapter 4, LUNC is permanently removed from the investable universe from May 2022 onward, so that a unique, idiosyncratic protocol failure does not dominate the comparative metrics that are meant to characterise allocation *rules*. To ensure that this intervention is not the source of the headline results, the entire scenario analysis is replicated under an *LUNC-unrestricted* variant in which the asset is allowed to remain in the eligible set after the collapse (i.e., the

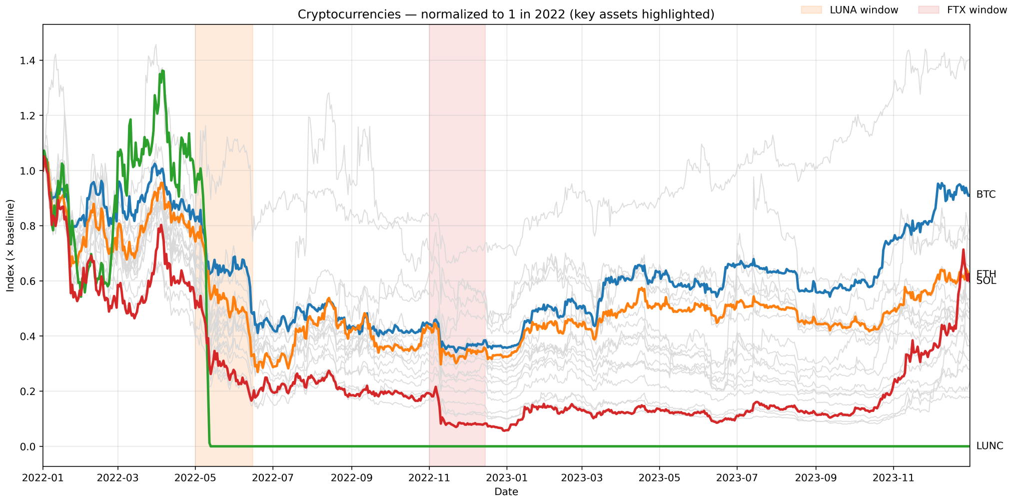

Appendix Figure B.1: Overview of the May 2022 Terra/LUNA collapse: per-strategy equity behaviour around the event window.

optimisers see its post-crash, near-zero price series). The associated panels are reported in Figure B.2 for the most exposed strategies (MSR, Return-Seeking utility, ERC and GMV), which compare the equity path and weight evolution with and without LUNC. The qualitative ranking of strategies is unchanged: GMV continues to dominate Scenario A on SR, ASR and CEQ, and MDC continues to dominate Scenario C, but the unrestricted variant amplifies the drawdown of return-seeking strategies without altering the regime-dependent conclusions.

The case study therefore serves three roles. First, it documents that a structural protocol failure inside the eligible universe is not simply "noise" but a directly observable amplifier of the estimation-error pathology of return-driven optimisers. Second, it justifies the single discretionary departure from the otherwise systematic selection rules of Chapter 3 by quantifying its effect on the headline metrics. Third, it links the empirical results to the broader institutional context of the 2022 crypto stress year, as further explored in Chapter 6 when transaction costs and execution feasibility are introduced.

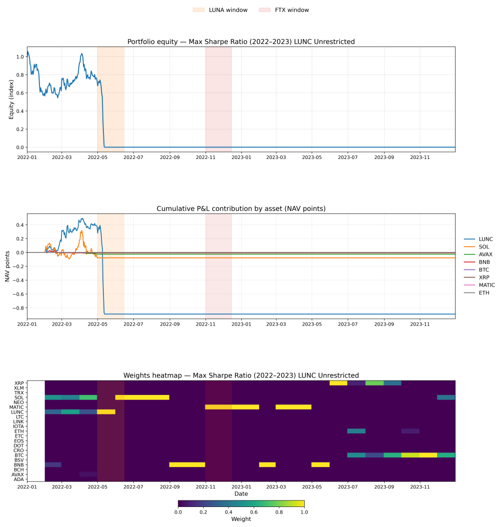

(a) MSR, LUNC-unrestricted

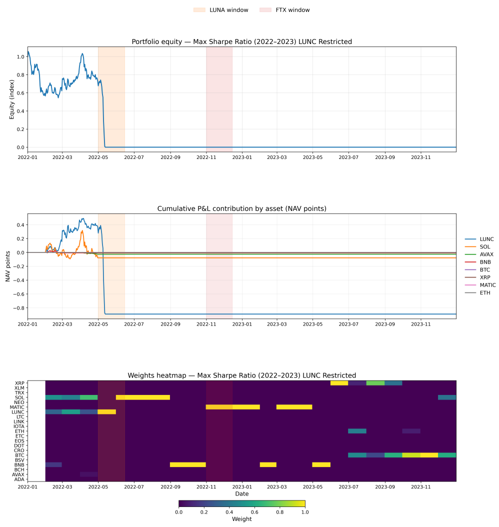

(b) MSR, LUNC-restricted

Appendix Figure B.2: Equity path and weight composition of the most LUNC-sensitive strategy (MSR) around the May 2022 collapse, with LUNC included in the eligible set (left) and removed after May 2022 (right). Analogous panels for Return-Seeking utility, ERC and GMV are reported in Appendix A.

Bibliography

Almgren, Robert and Neil Chriss (2001). "Optimal Execution of Portfolio Transactions". In: *Journal of Risk* 3.2, pp. 5–40.

Amenc, Noël, Felix Goltz, and Lionel Martellini (June 2013). *Smart Beta 2.0*. Tech. rep. EDHEC-Risk Institute.

Ammann, Manuel et al. (2022). *Survivorship and Delisting Bias in Cryptocurrency Markets*. SSRN Working Paper. DOI: 10.2139/ssrn.4287573. URL: https://ssrn.com/abstract=4287573.

Baltas, Nick and Robert Kosowski (2013). *Momentum Strategies in Futures Markets and Trend-Following Funds*. SSRN Working Paper. DOI: 10.2139/ssrn.1968996. URL: https://ssrn.com/abstract=1968996.

Binance (2024). *Trading Fees: Spot Trading Fee Schedule*. Binance. URL: https://www.binance.com/en/fee/schedule (visited on 03/15/2024).

Borri, Nicola (2019). "Conditional Tail-Risk in Cryptocurrency Markets". In: *Journal of Empirical Finance* 50, pp. 1–19. DOI: 10.1016/j.jempfin.2018.11.002.

Brière, Marie, Kim Oosterlinck, and Ariane Szafarz (2015). "Virtual Currency, Tangible Return: Portfolio Diversification with Bitcoin". In: *Journal of Asset Management* 16.6, pp. 365–373. DOI: 10.1057/jam.2015.5.

Briola, Antonio et al. (2022). *Anatomy of a Stablecoin's Failure: The Terra-Luna Case*. arXiv preprint arXiv:2207.13914. DOI: 10.48550/arXiv.2207.13914.

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay (1997). *The Econometrics of Financial Markets*. Princeton, NJ: Princeton University Press. ISBN: 9780691043012.

Caporin, Massimiliano et al. (2014). "A survey on the four families of performance measures". In: *Journal of Economic Surveys* 28.5, pp. 917–942.

Chan, Louis K. C. and Josef Lakonishok (1995). "The Behavior of Stock Prices Around Institutional Trades". In: *The Journal of Finance* 50.4, pp. 1147–1174. DOI: 10.1111/j.1540-6261.1995.tb04053.x.

Cho, Jaewoo (2023). "A Token Economics Explanation for the De-Pegging of the Algorithmic Stablecoin: Analysis of the Case of Terra". In: *Ledger* 8. DOI: 10.5195/ledger.2023.283.

Chopra, Vijay K. and William T. Ziemba (1993). “The Effect of Errors in Means, Variances, and Covariances on Optimal Portfolio Choice”. In: *The Journal of Portfolio Management* 19.2, pp. 6–11.

Choueifaty, Yves and Yves Coignard (2008). “Toward Maximum Diversification”. In: *The Journal of Portfolio Management* 35.1, pp. 40–51.

Christoffersen, Peter et al. (2010). “Is the Potential for International Diversification Disappearing?” In: *SSRN Electronic Journal*.

Clements, Ryan (2021). “Built to Fail: The Inherent Fragility of Algorithmic Stablecoins”. In: *Wake Forest Law Review Online* 11, p. 131. URL: http://www.wakeforestlawreview.com/2021/10/built-to-fail-the-inherent-fragility-of-algorithmic-stablecoins/.

CoinGecko (2025). CoinGecko. URL: https://www.coingecko.com (visited on 05/23/2025).

CoinGecko (Jan. 17, 2024). *2023 Annual Crypto Industry Report*. CoinGecko. URL: https://www.coingecko.com/research/publications/2023-annual-crypto-report (visited on 05/13/2026).

– (Jan. 15, 2026). *2025 Annual Crypto Industry Report*. CoinGecko. URL: https://www.coingecko.com/research/publications/2025-annual-crypto-report (visited on 05/13/2026).

Davis, M. H. A. and A. R. Norman (1990). “Portfolio Selection with Transaction Costs”. In: *Mathematics of Operations Research* 15.4, pp. 676–713. DOI: 10.1287/moor.15.4.676.

DeMiguel, Victor, Lorenzo Garlappi, and Raman Uppal (2009). “Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy?” In: *The Review of Financial Studies* 22.5, pp. 1915–1953.

Gârleanu, Nicolae and Lasse Heje Pedersen (2016). “Dynamic Portfolio Choice with Frictions”. In: *Journal of Economic Theory* 165, pp. 487–516. DOI: 10.1016/j.jet.2016.06.001.

Gensler, Gary (Jan. 10, 2024). *Statement on the Approval of Spot Bitcoin Exchange-Traded Products*. U.S. Securities and Exchange Commission. URL: https://www.sec.gov/newsroom/speeches-statements/gensler-statement-spot-bitcoin-011023 (visited on 05/05/2026).

Gkillas, Konstantinos and Paraskevi Katsiampa (2018). “An Application of Extreme Value Theory to Cryptocurrencies”. In: *Economics Letters* 164, pp. 109–111. DOI: 10.1016/j.econlet.2018.01.020.

Goltz, Felix and Sivagaminathan Sivasubramanian (May 2018). *Scientific Beta Maximum Decorrelation Indices*. Tech. rep. An ERI Scientific Beta Publication. ERI Scientific Beta.

Interactive Brokers (2024). *Stocks and ETFs Commissions: U.S. Tiered Pricing*. Interactive Brokers. URL: https://www.interactivebrokers.com/en/pricing/commissions-stocks.php (visited on 03/15/2024).

Khandoker, Mozaffar, Rafiqul Bhuyan, and Harminder Singh (2017). "Implementation Shortfall and Transaction Costs in Equity Markets". In: *International Journal of Finance and Banking Research* 3.6, pp. 91–98. DOI: 10.11648/j.ijfbr.20170306.12.

Kraft, Dieter (1988). *A Software Package for Sequential Quadratic Programming*. Technical Report DFVLR-FB 88-28. Oberpfaffenhofen, Germany: DFVLR (Deutsche Forschungs- und Versuchsanstalt für Luft- und Raumfahrt).

Leland, Hayne E. (1999). *Optimal Portfolio Management with Transaction Costs and Capital Gains Taxes*. Working Paper RPF-290. Research Program in Finance, Haas School of Business, University of California, Berkeley.

Liu, Yukun and Aleh Tsyvinski (2021a). "Risks and Returns of Cryptocurrency". In: *The Review of Financial Studies* 34.6, pp. 2689–2727. DOI: 10.1093/rfs/hhaa113.

– (2021b). "Risks and Returns of Cryptocurrency". In: *The Review of Financial Studies* 34.6, pp. 2689–2727. DOI: 10.1093/rfs/hhaa113.

Maillard, Sébastien, Thierry Roncalli, and Jérôme Teïletche (2010). "The Properties of Equally Weighted Risk Contribution Portfolios". In: *The Journal of Portfolio Management* 36.4, pp. 60–70.

Markowitz, Harry (1952). "Portfolio Selection". In: *The Journal of Finance* 7.1, pp. 77–91.

Michaud, Richard O (1989). "The Markowitz optimization enigma: is 'optimized' optimal?" In: *Financial Analysts Journal* 45.1, pp. 31–42.

Nansen (May 2022). *On-Chain Forensics: Demystifying TerraUSD De-peg*. Tech. rep. Nansen. URL: https://www.nansen.ai/research/on-chain-forensics-demystifying-terrausd-de-peg.

Osterrieder, Jörg and Julian Lorenz (2017). "A Statistical Risk Assessment of Bitcoin and Its Extreme Tail Behavior". In: *International Journal of Financial Engineering* 12.01, p. 1750038. DOI: 10.1142/S2010495217500038.

Petukhina, Alla et al. (2021). "Investing with Cryptocurrencies – Evaluating Their Potential for Portfolio Allocation Strategies". In: *Quantitative Finance* 21.11, pp. 1825–1853. DOI: 10.1080/14697688.2021.1880023.

Pezier, Jacques and Anthony White (2008). "The Relative Merits of Investable Hedge Fund Indices and of Funds of Hedge Funds in Optimal Passive Portfolios". In: *The Journal of Alternative Investments* 10.4, pp. 37–49.

Rodrigues, Daniel (2024). *Backtrader: Python Backtesting Library for Trading Strategies*. Backtrader. URL: https://www.backtrader.com/ (visited on 03/15/2024).

Rudin, Alexander M. and Jonathan S. Morgan (2006). “A Portfolio Diversification Index”. In: *The Journal of Portfolio Management* 32.2, pp. 81–89. DOI: 10.3905/jpm.2006.611807.

Sharpe, William F. (1964). “Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk”. In: *The Journal of Finance* 19.3, pp. 425–442.

The SciPy community (2026). *scipy.optimize.minimize: SLSQP method*. SciPy. URL: https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.minimize.html (visited on 04/30/2026).

U.S. Securities and Exchange Commission (Feb. 2023). *SEC Charges Terraform Labs and Do Hyeong Kwon with Defrauding Crypto Asset Investors*. Press Release No. 2023-32. U.S. Securities and Exchange Commission. URL: https://www.sec.gov/news/press-release/2023-32 (visited on 04/30/2026).

Uhlig, Harald (2022). *A Luna-tic Stablecoin Crash*. Working Paper 30256. National Bureau of Economic Research. DOI: 10.3386/w30256. URL: https://www.nber.org/papers/w30256.

Acknowledgments

I would like to thank my supervisor, Prof. Massimiliano Caporin, for his guidance during my studies and for the passion for the subject that he passed on to me, which encouraged me to propose and pursue this thesis.

I am deeply grateful to Elham, my rock, who stood by me at every step of this journey. Your help, patience, and support have meant more than I can adequately express.

I also wish to thank my family for their unwavering support, no matter what challenges arose along the way.

Finally, I'd like to thank my friends, Gianvittorio, Davide, Dani, Negin, Khashayar, Armin e Enrico, with whom I've shared my happiest moments and who have helped me become who I am today. And to all the people I've met throughout my life who, in one way or another, have shaped who I am: I'm grateful, even if I cannot name you all here.

Desidero ringraziare il mio relatore, Prof. Massimiliano Caporin, per la guida durante i miei studi e per la passione verso la materia che insegna, che mi ha incoraggiato a proporre e a realizzare questa tesi.

Sono profondamente grato a Elham, la mia roccia, che mi è stato accanto ad ogni passo di questo percorso. Il tuo aiuto, la tua pazienza e il tuo sostegno significano più di quanto io riesca ad esprimere.

Desidero inoltre ringraziare la mia famiglia per il loro sostegno incondizionato, indipendentemente dalle difficoltà incontrate.

Infine, vorrei ringraziare i miei amici, Gianvittorio, Davide, Dani, Negin, Khashayar, Armin e Enrico, con i quali ho condiviso i miei momenti più belli e che mi hanno aiutato a diventare chi sono oggi. E a tutte le persone che ho incontrato nella mia vita e che, in un modo o nell'altro, hanno plasmato chi sono: sono grato, anche se non posso nominarvi tutti qui.