A Multi-Factor Investment Framework for Crypto Asset Management

Alexander Aigner[†], Jérôme Benathan[*], Hugo Doering[*], and Dominik Stiftinger-Lang[†]

[*]Crypto United Partners Ltd., Baar, Switzerland
[†]ChainBrain FlexCo, Graz, Austria

Abstract

We present a comprehensive end-to-end investment framework for crypto asset management, covering capital inflows, investable universe definition, portfolio construction, proceeds allocation and rebalancing. The process integrates heterogeneous data sources and analytical dimensions (fundamental, quantitative, technical, and sentiment-based) into a unified and implementable architecture. We address the challenge to balance multiple valuation drivers and risk factors that differ in nature, magnitude and time horizon, evolving in an environment of fragmented and continuous news flow.

The investable universe is first defined through pre-specified evaluation criteria and screening filters, providing a disciplined foundation for any subsequent allocation. Fundamental assessments may, however, conflict with sentiment-driven narratives and behavioral biases, which can generate self-fulfilling price dynamics. Rather than treating these effects as residual noise, we model them explicitly through dedicated signals and views.

As crypto markets increasingly converge with traditional finance, feedback channels between the two ecosystems affect market momentum, regulation, and financial stability. To reconcile these dimensions, we employ an original portfolio construction engine that systematically optimizes, and validates asset exposures subject to risk parameters and investment guidelines. The resulting modular structure, with clear segregation of return sources, supports transparent performance attribution and governance. The methodology is suitable for institutional-grade deployment, while its individual modules can also be used as standalone strategies or as building blocks for digital asset products.

Keywords: Crypto asset management; Portfolio construction; Black–Litterman model; Multi-factor signals; Rebalancing execution; Investments

1 Introduction

The process of investing in traditional assets such as stocks and bonds is well established. The analysis of fundamental data serves as a safeguard during asset selection. Markowitz' initial contribution [Mar52] portfolio optimization has been an ongoing research topic of interest [Sal24].

For instance in the emerging and relatively recent digital asset class, some attention has been given to the effectiveness of Bitcoin as a hedge against inflation, with several studies evaluating its capacity to enhance the overall risk profile within diversified portfolios [Bor19; Gue+19]. Recent works also examine diversification benefits in pure cryptocurrency portfolios; with the exception of Brauneis & Mestel [BM19], however, only a limited investment universe of up to a dozen cryptocurrencies was used [JL17; Liu19; Ma+20]. Although cryptocurrency portfolio optimization research emphasizes the use of statistical characteristics (such as return distributions and volatility) to enhance performance, the actual selection of assets for the investment universe is often straightforward. In practice, it generally involves choosing leading cryptocurrencies by market capitalization, since these assets tend to offer superior liquidity and a longer trading history, traits critical for both analysis and efficient portfolio construction. Novel ideas for selection are to sort cryptocurrencies into sectors [ČTŽ22] and to use the number of tweets about individual cryptocurrencies [AMŠ21]. In the domain of cryptocurrency portfolio rebalancing, studies mainly

Published on December 17, 2025 in CVA Research Journal 2025.

Figure 1: **Investment process diagram**

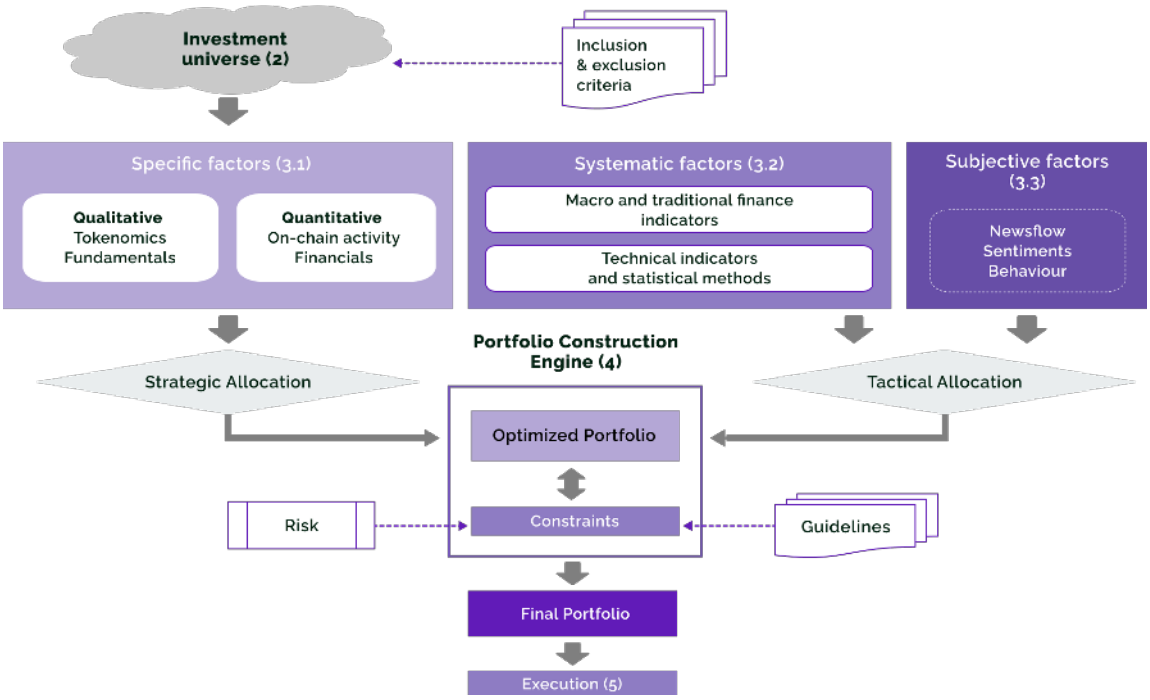

Note: The diagram summarizes the end-to-end workflow from capital provision to institutional implementation. Each numbered block corresponds to a section or subsection in the paper where that step is discussed in detail, including the investment universe and selection (Section 2), valuation driving factors (Section 3), portfolio construction (Section 4), and rebalancing and execution (Section 5).

address time-interval- and threshold-based rebalancing [HV21; SSN24], leaving out the challenges of actual rebalancing or reallocation execution.

This work aims to deliver a comprehensive, actionable framework for the entire cryptocurrency investment process from start to finish. Our goal is to move beyond fragmented and theoretical approaches by proposing investors and practitioners a systematic guide covering asset selection, risk controls, allocation strategies, and performance evaluation, so that decision making becomes both informed and effective for crypto portfolios. We develop an investment process that starts from a broad market-derived universe, applies economically motivated inclusion and exclusion criteria, aggregates heterogeneous valuation drivers (fundamental, quantitative and subjective) into a structured factor architecture, and feeds these inputs into a Black–Litterman-based portfolio construction engine. The resulting target allocations are implemented via a cost-optimal rebalancing and execution layer that explicitly accounts for trading costs and operational constraints. Empirically, we evaluate the framework on a multi-year sample of liquid cryptocurrencies and compare its performance to a range of market-capitalization and equal-weight benchmarks.

The contribution of this work is threefold: First, we provide an explicit, modular blueprint for an institutional-grade crypto investment process that is compatible with existing governance and risk-management practices. Second, we extend the Black–Litterman approach to a multi-factor setting tailored to digital assets, allowing heterogeneous signals and subjective views to be combined in a controlled way, and provide a smart solution to reallocate crypto portfolios. Third, we document the properties of the resulting strategy, including its behavior across different market regimes.

Figure 1 summarizes the overall workflow of the proposed framework, from the definition of the investable universe through factor aggregation and portfolio construction to rebalancing and operational deployment. The subsequent subsections briefly interpret this diagram in the broader context of institutional portfolio management and position our study within the investment life cycle.

The paper is structured as follows: Section 2 (Investment Universe and Selection) introduces the investment universe and the inclusion and exclusion criteria used to define the set of eligible assets. Section 3

Figure 2: **Investment life cycle for professional portfolio management**

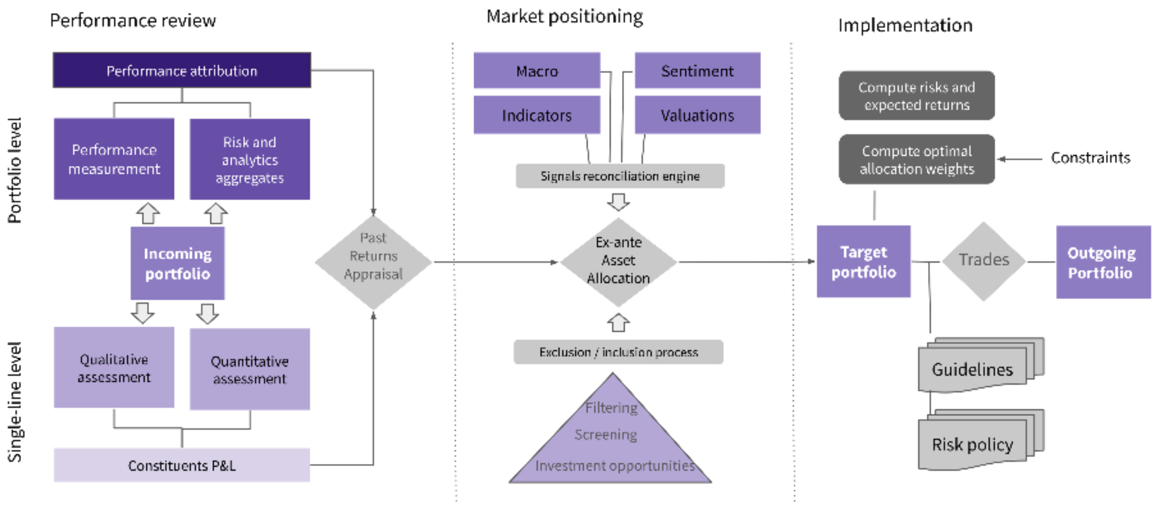

Note: This diagram provides a holistic overview of how professional portfolio management teams dynamically allocate assets, appraise performance, and iteratively rebalance investment portfolios in response to changing market conditions and governance frameworks.

(Valuation Driving Factors) describes the valuation driving factors and factor families that generate signals for allocation decisions. Section 4 (Portfolio Construction Engine) presents the portfolio construction engine, including data integration, the Black–Litterman framework, the constrained optimization problem, and the empirical performance results. Section 5 (Rebalancing Execution) discusses the rebalancing logic, trading-cost modeling, and execution considerations. We then summarize the broader implications in Sections 6 (Discussion) and 7 (Conclusion).

Performance Analysis and Investment Life Cycle

While Figure 1 focuses on the internal mechanics of our framework, institutional investors typically embed any strategy into a broader investment life cycle that combines performance review, market positioning and implementation. Figure 2 provides a complementary view of this cycle and clarifies how the components developed in this paper can be used as a guideline for professional investing in digital assets.

The cycle begins with a comprehensive evaluation of the incoming portfolio, where the performance is dissected at both the position and aggregate levels. Individual holdings undergo qualitative and quantitative assessment, examining each constituent's profit and loss contribution, while portfolio-level analytics measure overall performance and risk exposures. These inputs feed a performance attribution analysis (allocation, security and interaction effects) that reveals which investment decisions added or detracted value and forms the basis for subsequent positioning.

On this basis, portfolio managers move to forward-looking market positioning. Multiple information streams such as macro conditions, technical indicators, market sentiment and fundamental valuations are consolidated into coherent investment views via a signals reconciliation engine. The resulting ex-ante asset allocation targets are then filtered through inclusion and exclusion criteria, screening methodologies and opportunity identification, narrowing the investment universe to securities that align with both quantitative signals and qualitative judgment.

In the final stage, target allocations are transformed into actual portfolio positions. An optimization engine computes allocation weights subject to regulatory limits, risk budgets, liquidity requirements and investment guidelines. The resulting target portfolio serves as the blueprint for trading operations that transition the existing holdings into the desired configuration. This process is iterative: the outgoing portfolio from one cycle becomes, after re-allocation and subsequent trades execution, the incoming portfolio for the next sequence, enabling ongoing performance assessment and strategy refinement.

Figure 3: **Relative market share among the top 10 altcoins**

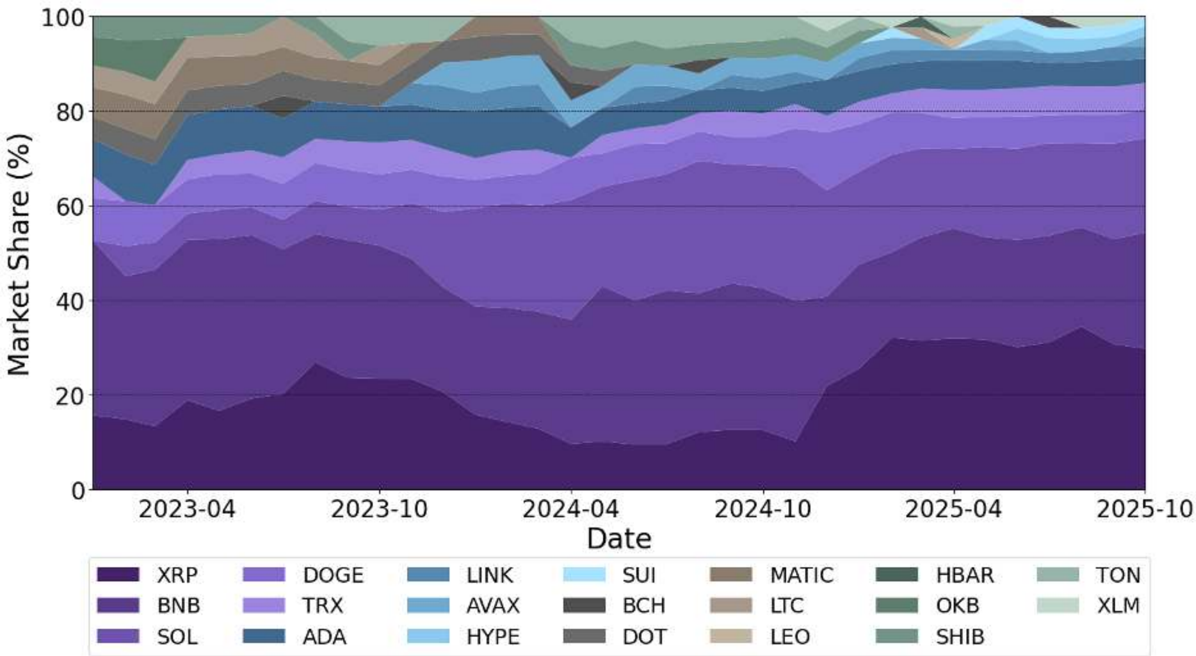

Note: The top cryptocurrencies excluding Bitcoin, Ethereum and stablecoins were picked on the first day of each month in the depicted period.

2 Investment Universe and Selection

2.1 Definition of the Investment Universe

An investment universe in crypto refers to a complete set of digital assets and tokens that are eligible for selection, analysis, and potential inclusion in a portfolio based on specific investment criteria. This universe is established by defining rules and parameters such as market capitalization, liquidity, technology category, and sector inclusion or exclusion. In practice, the investment universe may contain large, well-established cryptocurrencies (like Bitcoin or Ethereum), promising existing projects with technological innovation, and, depending on strategy, a limited allocation to emerging tokens with high appreciation potential. Institutional investors often further narrow the universe to sectors such as Layer 1 and Layer 2 blockchains or decentralized finance, while explicitly excluding areas like stable coins, metaverse, gaming, or memecoins based on qualitative and quantitative selection as well as a set of ex-ante exclusion criteria. The definition and ongoing refinement of the investment universe is a critical "first line of defense" in a systematic and disciplined crypto asset management process. In the context of a fat-tailed distribution of investment opportunities combined with a rapidly evolving crypto investment landscape, the investment universe must be circumscribed. This voluntary restriction allows for a manageable portfolio where each exposure can be effectively monitored, regularly reviewed and reassessed for ongoing suitability and eligibility in the investment strategy.

2.2 Systematic Implementation Approaches

A predefined investment universe serves as a structured foundation, allowing the asset manager to construct a portfolio either by strictly selecting exposures from within this set or by using it as a reference framework from which deviations ("off-benchmark" positions) can be made to express high-conviction views.

Index-based approach. The digital asset universe consists of thousands of digital assets, each with unique characteristics, risks, and rewards. Indices are hence practical investment aggregates, as they conduct both a qualitative and quantitative evaluation of the price dynamics and define a representative typology of the current crypto currency market. The underlying components have to meet eligibility criteria screened on their liquidity, tradability on vetted exchanges, and availability with vetted custodians.

Table 1: **Asset exclusion criteria per sector**

Sector	Exclusion Criteria
Layer 1 blockchains	Assets are excluded if they lack core infrastructure functions such as smart contract capability, cross-chain compatibility, or decentralized liquidity. Projects with limited development activity or without a verifiable consensus mechanism (e.g., Bitcoin forks or inactive chains).
Layer 2 ecosystem	Protocols that are not true scaling solutions (i.e., do not handle off-chain computation or transaction aggregation) or rely entirely on Layer 1 mining and validation.
DeFi assets	Centralized or hybrid protocols lacking permissionless access, showing high counterparty risk, or deviating from open-source governance. Stablecoins and privacy-centric tokens are excluded due to low decentralization or compliance concerns.
AI crypto sector	Tokens whose core use cases overlap excessively with unrelated industries (e.g., data privacy or cloud computing) or face unresolved regulatory issues around data handling or securities classification.
Metaverse and gaming tokens	Excluded for weak legal frameworks around ownership/IP. Projects that issue NFTs without enforceable ownership rights or rely on centralized control mechanisms are disregarded.

Different classic methods can be applied to the index construction. However, the diversity of protocols and tokenomics complicates research.

Some index providers may use different proprietary liquidity weightings methodology or add relevance and innovation scores. Some indices, particularly those focused on specific sectors like DeFi, include cryptocurrencies based on their role and influence within the crypto ecosystem. For example, tokens driving significant innovation in decentralized finance might be selected for a DeFi-focused index. This may address the fact that index weightings often may not reflect the relative prominence of the projects.

Rule-based approach. If "hard plugged" quantitative criteria may offer less room for interpretation, blockchain protocols vary considerably in their technical architecture and operational mechanisms. Consequently, the metrics that are meaningful for one protocol may be irrelevant, non-existent, or incomparable across others, rendering data aggregation and cross-protocol analysis inherently complex or inaccurate. Similarly, in traditional finance, the debt capacity and leverage of an investment-grade cyclical multinational corporation may differ substantially from a regulated domestic utility, even though both may be classified as "corporate debt" embedding default risk. Moreover, as can be seen in Figure 3, market capitalization, the most commonly employed metric in equity markets, may prove inadequate in the context of cryptoassets, where pronounced valuation fluctuations would inevitably result in significant and frequent index rebalancing and generate periodic portfolio trades.

As mentioned in Bartolucci and Kirilenko [BK20] "cryptoassets can be described according to two features: security (technological) and stability (governance). Investors' perspective is being driven by (i) their attitudes towards assets' features, (ii) information about the adoption trends, and (iii) expected future economic benefits of adoption." Exclusion criteria across all sectors often include assets with opaque governance, insufficient liquidity, unreliable pricing data, or lack of secure custody solutions. The sector-based filtering approach, presented in Table 1, ensures that only transparent, technically sound, and institutionally viable assets are eligible for sector-based crypto investment frameworks such as the FTSE Grayscale indices [FTS] or the S&P Global Top 15 DeFi index [SP].Designing inclusion and exclusion criteria leaves ample room for sophistication and can be translated into measurable factors and related data-driven approaches.

3 Valuation Driving Factors

Valuing cryptoassets requires departing from traditional valuation methods (i.e., discounted cash-flow-based models, EV/EBITDA) to account for multiple interacting drivers of risk and return. This section takes a holistic view of these drivers, from ratings and qualitative assessments to structural features and market sentiment.

3.1 Ratings and Qualitative Factors

In traditional finance, ratings describe the quality of an existing or prospective investment. As a normative tool they ensure comparability in a heterogeneous universe. In fixed income, credit ratings help investors assess an issuer's solvency and manage exposure to default risk.

The main strength of ratings also constitutes their main weaknesses. On the one hand, they must offer an assessment with some "staying power" as the current grade and its outlook has to factor in any forecasts and future elements that could affect the financial profile. On the other hand, while expressing a static opinion at a given moment in time, their stability could forego some short-term events and lag any rapidly evolving situation which could in turn lead to multi-notch downgrades or upgrades. In addition, grades or ratings are not a market measure of risk such as bonds' credit spreads, option adjusted spreads or credit default swaps which might adjust in real time to an evolving news flow or investor's perception of the creditworthiness of an issuer.

The "balancing act" in fixed income investing aims to reconcile two analytical dimensions: credit ratings expressing a likelihood of default over a given period of time derived from historical data, and credit spreads translating in real-time market-implied default probabilities, thus enabling a direct comparison between ratings and observed market spreads. How can this well-established practice in traditional finance be adapted to cryptocurrencies and digital assets, providing a broader framework for evaluating risk, valuation, and market efficiency within decentralized financial ecosystems?

3.1.1 Translation of a Normative Approach to Crypto

Widely recognized agencies such as Moody's, S&P or Fitch do not at this stage systematically assign grades to crypto or have established clear "investment-grade", "speculative or high-yield" thresholds, nor do token issuers apply for a public and independent rating each time they envisage to issue tokens or plan an ICO. However, several initiatives from crypto research houses, index providers or product originators mark the beginning of normative approaches and some uniformization[1], attempting to synthesize multiple risk factors into a single, easily interpretable metric. Some simple but effective methodologies focus on five dimensions: governance, integration, liquidity, security and operational efficiency each subdivided into more granular subcriteria[2] and designed to be "use neutral" across sectors such as Layer 1 protocols, exchanges or DeFi platforms. Figure 4 illustrates a stylized scoring framework of this type, in which each dimension is rated on a fixed 1–5 scale and combined into a single summary score using ex-ante weights. This approach is reminiscent of early credit-scoring models such as the Altman Z-score [Alt68] and allows large universes of assets to be screened and ranked with a simple transversal method, at the cost of potentially overlooking idiosyncratic features of individual assets.

The introduction of standardized crypto ratings[3] undoubtedly enhanced transparency and legitimacy upon protocols or token issuers by aligning them with established evaluation methodologies. A rating framework facilitates engagement by regulated financial entities such as banks and investment management companies by providing a familiar "bucket-based" risk assessment standard. Additionally, it supports more systematic benchmarking, enabling market participants to determineexpected returns based on standardized risk metrics rather than protocol-specific factors alone.

[1] The traditional credit rating scale from both agencies ranges from AAA/Aaa (highest creditworthiness) through investment-grade categories (AA/Aa, A, BBB/Baa) to speculative-grade or "junk" categories (BB/Ba, B, CCC/Caa, and below). This classification creates a critical threshold at the investment-grade boundary, as many institutional investors face regulatory or internal policy constraints that restrict holdings to investment-grade securities

[2] For example, financial transparency can include audited financial statements, proof-of-reserves and other disclosures supporting claimed stability and solvency.

[3] On 8 August 2025, S&P Global Ratings assigned its first credit rating to a DeFi protocol, evaluating Sky Protocol (formerly MakerDAO) at B- with a stable outlook on selected stablecoin-related liabilities. The rating does not cover governance tokens and should be viewed as one component within a broader risk framework.

Figure 4: **Selection and evaluation process example**

Category	Dimension	Question	
MODEL:	CONCEPT:	Is the proof of concept coherent and understandable?	1
	REAL NEED:	Does it address a real problem or need?	2
	VALUE PROPOSITION:	Is there a clear value proposition for all groups to participate?	3
	VISION:	Is there a vision beyond the immediate concept that matches broader trends?	4
TECHNOLOGY:	FEASIBILITY:	Does the concept look like it could work?	2
	MVP:	Has it been tested?	3
	DECENTRALIZED:	Is it already or is there a roadmap in place to become truly decentralized?	4
	INNOVATION:	Is the product innovative? Does it contribute to the blockchain ecosystem?	4
TEAM:	FOUNDERS:	Are the founders known & do they have any relevant experience?	1
	ADVISORS:	Do they have reputable advisors with relevant experience?	1
	VESTING:	Does the team have a sufficient stake & vesting for balanced incentives?	3
	CONNECTED:	Are the founders well-connected in the ecosystem?	4
	SKILLS:	Do founders & advisors cover all areas (business, tech, marketing, industry)?	4
MARKET:	MARKET POTENTIAL:	Is there a large market & growth potential?	3
	COMPETITION:	Are they outstanding in their field vs. competitors, if any?	3
	DEFENSIBILITY:	Can the team defend their business against others (IP, network effects, brand, ..)?	4
ECONOMICS:	NEEDS BLOCKCHAIN:	Does the project benefit from using blockchain technology?	1
	TOKEN UTILITY:	Does the ecosystem get added and/or inherent value from having a token?	1
	TOKEN INCENTIVES:	Do the token economics create growth incentives & network effects?	3
	TOKEN ECONOMY:	Are there incentives to grow beyond the limits of a medium of exchange?	4
TRACTION:	HYPE & MEDIA:	Is there positive interest on (social) media and chats?	1
	ROADMAP:	Do they have an ambitious yet realistic roadmap?	2
	PARTNERSHIPS:	Are there already high-profile partners working with them?	3
	PROGRESS:	Has an alpha or beta already launched?	4
	CUSTOMERS:	Is there an existing customer base?	4
FINANCIALS:	VALUATION:	Is the valuation reasonable?	1
	RUNWAY:	Will the funding be sufficient but not too high for the scope of the project?	2
	SUPPLY SOLD:	Does the team distribute a reasonable amount of tokens?	2
	DISCOUNTS:	Are discounts reasonable in relation to risk profile of rounds?	3
	EXCHANGES:	Are there already exchange listings in preparation?	4
		Score	81

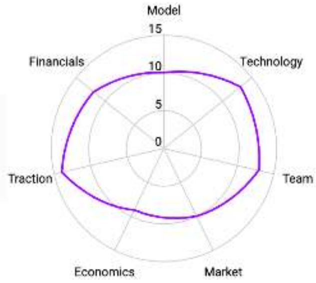

Rating
AAA
AA+, AA, AA-
A+, A, A-
BBB+, BBB, BBB-
BB+, BB, BB-
B+, B, B-
CCC+, CCC, CCC-
CC
C
D

Note: Methodology inspired by Boris Golden's "4 Ms" for early-stage startups [Gol]. Each dimension is rated on a 1–5 scale; the overall score is the average across dimensions. Scale meanings: 1 = No information / suspected scam; 2 = Clear weakness / risky; 3 = Neutral; 4 = Convincing; 5 = Outstanding / exceptional.

Conversely, the rating's scope can be narrowly restricted to certain idiosyncratic properties, failing to accurately encompass the broader risk landscape of a specific protocol. Ratings themselves represent static snapshots and are often unable to adapt to the rapidly evolving nature of operational risks inherent within blockchain-related entities or decentralized finance. The application of legacy financial risk models may inadequately reflect blockchain-native threats or inherent risks, and there is potential for non-specialist investors to misinterpret the rating as an unqualified indication of safety or price stability.

Furthermore, cryptoassets, particularly unbacked cryptocurrencies, have no issuer obligation to repay principal or make periodic payments. Instead, crypto-related scores attempt to evaluate a combination of technological viability, network adoption, and market dynamics. This could create conceptual confusion. A cryptoasset with excellent technology and strong adoption can experience severe price declines due to market sentiment, regulatory actions, or other macroeconomic factors, none of which indicate an intrinsic deterioration. On the contrary, a technologically inferior asset may maintain value through first-mover advantage or network effects, or experience a material price increase due to orchestrated promotion via influencers or counterparties supporting the price via market activity[4].

As alternative to these "public rating" which are mostly unsolicited and freely available, crypto audit scoring, particularly for smart contracts and token issuers, is generally a paid service provided by an auditing firm.[5] Focusing on similar criteria as ratings, this process systematically examines a protocol's smart contract code, governance mechanisms, team transparency, incident history, and mitigation of vulnerabilities, culminating in a numeric or letter-based score that reflects overall risk. Applied to cryptoasset management, these scores could be integrated into due diligence frameworks, exclusion criteria, or quantitative models to screen investments.

[4] Cryptocurrency markets suffer from rampant wash trading, with studies estimating that over 70-95% of reported trading volume on unregulated exchanges represents artificial activity. Exchanges engage in wash trading to inflate rankings on data aggregators like CoinMarketCap, creating false signals of liquidity and adoption. A 70% wash trading ratio can move an exchange's rank up by 46 positions. This data pollution undermines any rating methodology that relies on trading volumes, market capitalization, or liquidity metrics as inputs[Con+23].

[5] See major crypto audit providers such as CertiK and Hacken. Switzerland-based audit firms such as ChainSecurity and ConsenSys also stand out for local expertise and global reach.

Table 2: **Basel prudential standards for cryptoassets**

Basel Group / Tier	Description & Examples	Eligibility & Risk Criteria	Regulatory / Capital Treatment	Portfolio Allocation Guidance
Group 1a (Tier 1)	Regulated, tokenized traditional assets	Meets all "Group 1" classification conditions: tokenized securities, fully regulated, redeemable, robust legal structure, transparent audit, effective governance.	Same capital rules as the underlying asset; eligible collateral (subject to assessment); infrastructure risk add-on possible.	Allocate akin to underlying (e.g., cash equivalents / short bonds) subject to ratings; apply standard concentration limits.
Group 1b (Tier 1)	Regulated (non-algorithmic) stablecoins	Collateralized, issued by supervised entity with transparent redemption, regular audits, strong risk management; meets regulatory tests.	Dedicated capital charge (Pillar 1); strict redemption-risk management and reserve audits; infrastructure risk add-on possible.	Potentially treated like short-duration liquid assets, per collateral quality; typically $<$ 10% of portfolio.
Group 2a (Tier 2)	Large-cap unbacked crypto (BTC, ETH, etc.)	Fails some Group 1 criteria but meets hedging-recognition standards (market cap, liquidity, regulatory clarity).	Modified standardized approach; not eligible for internal models; exposure cap $<$ 2% of Tier 1 capital.	Consider 1–10% tactical allocations within a research-backed, dynamically rebalanced framework.
Group 2b (Tier 3)	Emerging protocols / altcoins	Fail Group 1 and most Group 2a criteria; niche tech, thin markets, incomplete regulatory/operational standards.	Conservative treatment: 1250% risk weight; long and short exposures aggregate for risk.	Small satellite allocations (0.5–2%) only after rigorous due diligence; not suitable as core holdings.
Group 2b (Tier 4)	Highly speculative / non-compliant assets	New tokens with weak transparency, governance or audits; significant DeFi governance risks; no regulatory recognition.	No collateral recognition; extreme capital charges; disclosure required; exposure discouraged or prohibited.	Excluded from institutional portfolios regardless of private ratings.

Source: Bank for International Settlements [21]. **Note:** The finalized Basel Cryptoassets Standard is a significant step toward preparing global financial systems for the wide adoption of distributed ledger technology (DLT). The revisions demonstrate responsiveness to stakeholder calls for flexibility and discretion instead of rigid rules before risks have fully materialized. While the Basel Committee remains very conservative about cryptoasset risks, especially regarding public blockchains, the standard better aligns with the market view that some regulated cryptoassets (e.g., tokenized securities, quality stablecoins) should receive risk treatment similar to conventional assets when they meet strict regulatory and operational criteria. Crucially, the Standard only sets minimum requirements—national authorities may adopt stricter rules, especially after recent failures of crypto intermediaries and intense scrutiny by policymakers. Even before full implementation and adoption, the Basel rules are likely to shape bank decision-making, risk appetite, and global regulatory harmonization for cryptoasset exposures.

3.1.2 Alternative Framework: Risk-Based Classification vs. Creditworthiness Ratings

Rather than replicating the normative regulatory function of credit or equity ratings, crypto ratings or equivalent assessment such as the framework proposed by the BIS (Bank for International Settlements) are most useful as risk classification tools. As an example, Basel Committee's Group 1/Group 2 framework in Table 2 provides a precedent: "Group 1" assets must satisfy strict eligibility criteria (regulatory compliance, audits, credible stabilization mechanisms), whereas "Group 2" assets attract substantially higher capital charges and are treated as speculative exposures.

This offers a practical classification scheme and regulatory based on observable attributes and safeguards rather than subjective rating judgments. It is also compatible with behavioral portfolio theory, which emphasizes that investors segment portfolios by goals and risk tolerance instead of pursuing a single mechanically "optimal" allocation.

3.2 Macroeconomics and Systemic Factors

Macroeconomics and cryptoassets have traditionally been viewed more as distant relatives than as closely connected twins in terms of interaction or any form of mutual influence between the two ecosystems. This perception largely stems from the relatively modest size of the emerging crypto asset class compared to the vast capitalization of secular traditional markets. In addition, the existence of two distinct universes, each with its own dedicated investors and communities, focused primarily on their respective performance, with a low degree of attention or interest in developments in the other.

Within the scope of our investment process, we systematically reviewed the empirical literature for each relevant period and endeavored to quantify the evolving dynamics of these complex relationships.

3.2.1 Literature and Research Review

Macro news. Corbet et al. [Cor+20], investigate how major macroeconomic news announcements influence Bitcoin returns through a constructed sentiment index derived from relevant news coverage. Unlike equities, where positive economic releases about unemployment rates and durable goods orders or consumer confidence for instance typically leads to price gains, the study finds that improvements in these indicators actually correspond with decreases in Bitcoin returns. Conversely, negative news on unemployment and durable goods is associated with increases in Bitcoin returns. The paper observes no statistically significant relationship between Bitcoin returns and news related to GDP or the Consumer Price Index (CPI). These findings indicate that the Bitcoin dynamics behave differently in response to mainstream macroeconomic news that traditional assets, underlining Bitcoin's distinctive role in financial markets.

The PPI.[6] As noted by Scharnowski [Sch24], in the face of economic uncertainty or inflationary pressures, investors may be more inclined to choose traditional safe-haven assets such as gold or government bonds, rather than Bitcoin, due to its high price volatility and lack of long-term value storage function like gold. However, if traditional assets perform poorly during such periods of economic stress, investors seeking high-risk, high-return investment opportunities may paradoxically turn to crypto as an alternative source of return. The same logic may apply to a weakening dollar index that might present an inverse correlation with Bitcoin, albeit the relation USD/BTC may not be uniform and always significant [LJG24].

Treasury yields mirror shifts in the macroeconomic environment and signal investor sentiment. According to Abakah et al. [Aba+20] treasury yields should be negatively correlated with cryptocurrency returns. Intuitively, when yields rise, they often reflect expectations for stronger growth or higher inflation, prompting investors to prefer traditional high beta assets presenting "superior quality returns" and reduce exposure to riskier ones like Bitcoin. Liu et al. [LJG24] observe a more complex and less straightforward relationship between concluding that the positive impact of bond yields on Bitcoin returns reveals that investors consider fixed income returns as a positive indicator of market sentiment and risk appetite. In other words, according to Liu et al. [LJG24], when government bond yields increase, it could present a favorable opportunity to enter the Bitcoin market, which contradicts the results from Abakah et al. [Aba+20] conducted five years earlier.

Monetary policies and macro news announcements (Part 1). Enlarging the area of investigations, Benigno and Rosa [BR23] investigates whether Bitcoin prices react to macroeconomic and monetary policy news by leveraging high-frequency intraday data (5 minutes frequency for Bitcoin and 30 minutes window around major U.S. macroeconomic and monetary policy announcements)[7]. The authors found out that, unlike other traditional U.S. asset classes, traditional assets (stocks, gold, major currencies), Bitcoin returns show negligible and statistically insignificant reactions to virtually all forms of monetary surprises and most macroeconomic news, with only a weak link to selected CPI releases. The authors modeled Bitcoin as an asset with no intrinsic value for which its current price depends on the discounted

[6] The Producer Price Index (PPI) is an important indicator for measuring the health of the macroeconomic environment. It reflects the price changes of various raw materials, fuels, energy, and other production materials required in the production process

[7] Benigno and Rosa [BR23] collected different macro news: news on the real economy (payroll, initial jobless claim, industrial production, retail sales, unemployment, and trade balance); news on inflation (PPI and CPI); and news on forward-looking indicators (confidence and ISM manufacturing), where the Bloomberg consensus forecast serves as a proxy for market expectations. Shocks are disentangled into: Target rate changes (unexpected moves to the policy rate), Path factor (unanticipated future rate guidance), and LSAP surprise (large asset purchases). These factors are extracted using principal components from relevant futures markets.

Figure 5: **Bitcoin price impulse responses by risk regime and shock**

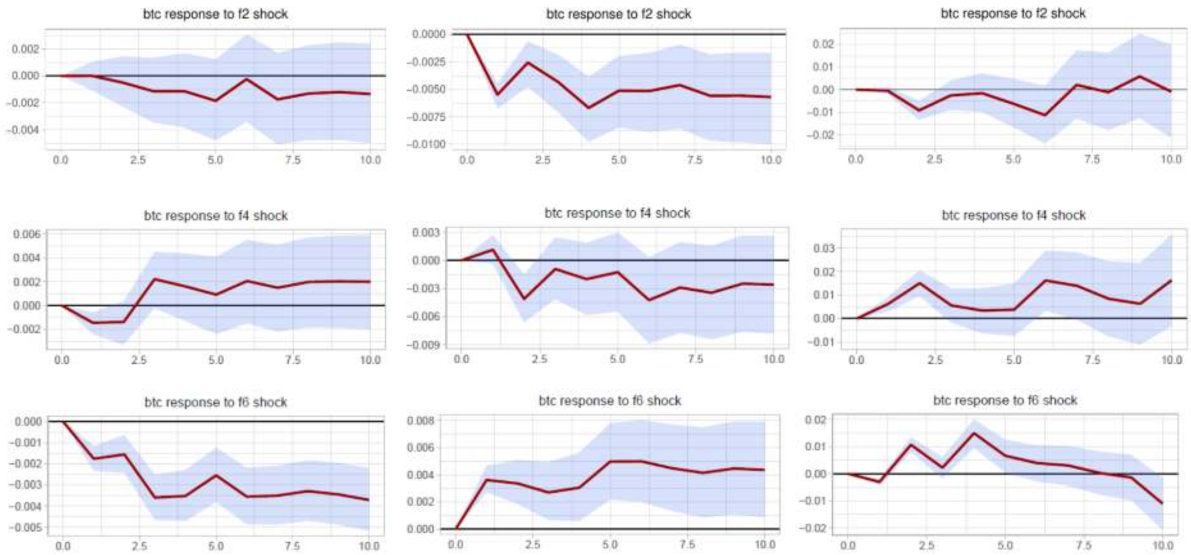

(a) Low risks: $\Delta f_{1,t-1} < 3\%$ (b) Medium risks (c) High risks: $\Delta f_{1,t-1} > 10\%$

Source: Thélissaint [Thé25] **Note:** Bitcoin relative price movements for 10 days after different macro shocks (f2), financial stress (f4) and market optimism (f6) under various risk regimes (low risks, medium risks, high risks). For details see [Thé25].

value of its future price. Interestingly, the result that Bitcoin does not react to monetary news is rather surprising as it casts some doubts on the role of discount rates in pricing Bitcoin.

Monetary policies and macro news announcements (Part 2). With similar intentions Mužić and Gržeta [MG22] tested the main hypothesis that Bitcoin is affected by macroeconomic news expectations on gold, the S&P 500, 2-year Treasury bills, and 10-year Treasury bills in the same way as other traditional assets. Applying sophisticated GARCH and ARMA models[8], the authors however conclude that U.S. macroeconomic news has a measurable effect on Bitcoin returns (12 significant coefficients), but limited influence on trading volume, and does not significantly drive volatility. Notably, Bitcoin's return sensitivities more closely align with 2-year Treasury notes than gold, indicating that Bitcoin cannot be regarded as "digital gold" from a macroeconomic perspective. Surprisingly, the S&P 500 had the lowest response. The evidence supports the view that Bitcoin responds to macroeconomic expectations similarly to other risk assets, though with inconsistent directions. They conclude that Bitcoin remains dominated by news-driven sentiment rather than fundamental valuation.

An integrated framework. In a very comprehensive study Thélissaint [Thé25] uses a Factor-Augmented Vector Autoregression (FAVAR) model to identify six key risk factors driving cryptocurrency markets, revealing three distinct market regimes: low-, medium-, and high-risk. In low-risk times, volatility risk dominates returns; medium-risk environments are influenced by macro-confidence; while high-risk periods see DeFi Risk-On behavior, optimism, and equity reallocation jointly shaping market direction. Strengthened macroeconomic confidence and financial stability typically benefit crypto markets, though returns may decline due to reduced risk premiums. As highlighted in Figure 5 and noted by the author, cryptocurrency prices (as represented by Bitcoin in this case) do not respond uniformly across risk regimes, both in terms of magnitude and response velocity. In stressed markets, however, heightened volatility risk caused by falling macro-confidence triggers a flight to safety, reducing capital inflows and optimism in crypto, shifting cryptocurrencies as complementary assets substitutes such as equities. The study also finds that Bitcoin prices can feedback into financial stress and macro-confidence, indicating growing spillover effects from crypto to the broader real economy.

[8] GARCH (Generalized Autoregressive Conditional Heteroskedasticity) models are appropriate when volatility is autocorrelated and exhibits mean reversion, such as during periods of market turbulence or crises. ARMA models (autoregressive Moving Average models) models were used for apprehending to account for "volume clustering".

Figure 6: **Rolling correlations vs Bitcoin**

Note: The rolling correlation of the Bitcoin price to S&P 500 (SPY), Gold price, US 10 years Treasury Yield (US10Y), CBOE Volatility Index (VIX), and the Cryptocurrency Index 30 (CCI30, Top 30 excl. stablecoins)

3.2.2 Implications for Portfolio Management and Practitioners

As observed, the complex relationship between cryptocurrencies and macroeconomic factors has attracted considerable attention. Yet, despite advanced models and sophisticated statistical methods, no consistent or replicable pattern has emerged. While impulses from the real economy are not entirely disconnected from cryptoassets, the primary valuation drivers appear to be conveyed through subjective channels such as risk aversion or broad sentiment toward speculative assets, rather than through mechanical transmission mechanisms or a standardized cryptoasset pricing framework. Moreover, joint dynamics, self-reinforcing sentiment, feedback loops, and shifting attention patterns make the isolation of a single or even a dominant set of valuation drivers especially challenging.

In this context, the growing interactions between cryptocurrencies and traditional markets carry significant implications for portfolio diversification. As institutional participation deepens and macro-financial linkages intensify, allocation frameworks will need to integrate both traditional and crypto dynamics within a single risk perspective. One practical way to monitor these evolving linkages is to track rolling correlations between cryptoassets and major equity, rates, and volatility benchmarks. Over the past year, the 60-day rolling correlation between the SPY and Bitcoin has fluctuated between 0.25 and 0.5[9].

Figure 6 displays the 60-day rolling correlations between Bitcoin (BTC) and several major financial market instruments and indices: the CCI30 Crypto Index, gold, the S&P 500 (SPY), U.S. 10-Year Treasury yields (US10Y), and the VIX volatility index. It tracks how these relationships evolve over time between March 2024 and November 2025. We can observe the following:

- The correlation between BTC and the CCI30 Crypto Index remains persistently high and positive, reflecting Bitcoin's strong integration with the broader crypto market.
- Bitcoin's correlation with SPY (S&P 500) trends upward, suggesting stronger alignment with U.S. equities through mid-to-late 2025.
- Correlations with gold and the 10-year Treasury yield (US10Y) fluctuate near zero, indicating that Bitcoin does not consistently behave like either gold or bonds.
- The VIX index shows a negative correlation with BTC, especially pronounced at certain intervals, implying BTC sometimes moves opposite to spikes in market fear.

[9]The Pearson correlation coefficient quantifies the degree of linear association between two variables. A correlation approaching +1 indicates strong co-movement, whereas a value near 0 suggests independence.

Figure 7: **Rolling correlations vs NASDAQ and Correlation among major crypto currencies**

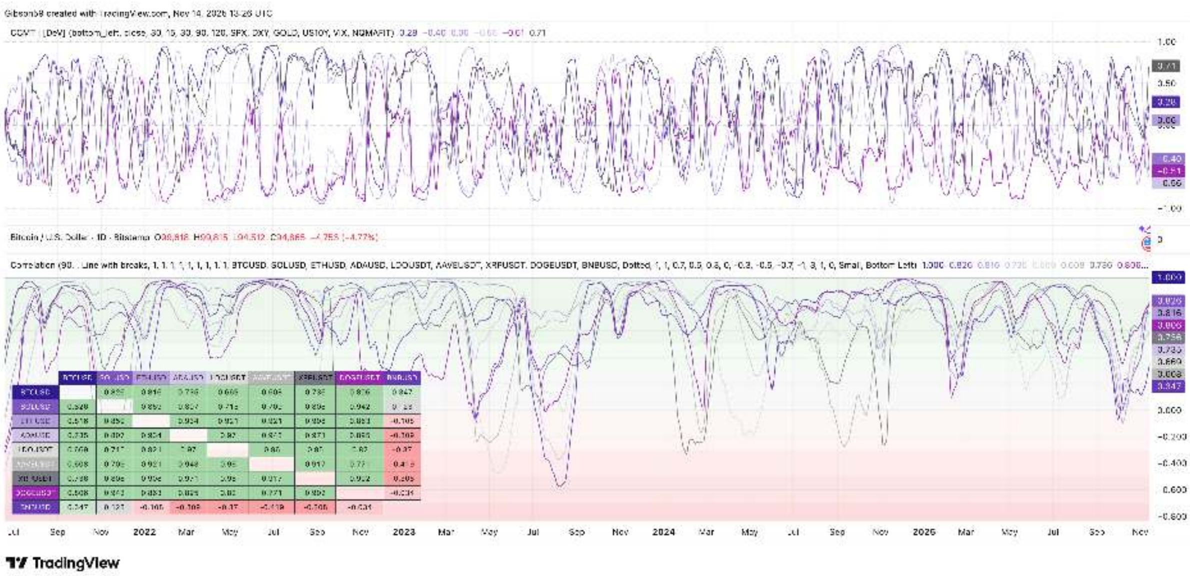

Source: Trading View. **Note:** The upper chart tracks the correlation coefficients between NASDAQ (NQMAFIT) and key macro financial indicators: S&P 500 (SPX), US Dollar Index (DXY), Gold, US 10-Year Treasury Yield (US10Y), and VIX. The lower chart highlights inter-crypto correlations among major coins like BTC, SOL, ETH, ADA, LDO, AAVE, XRP, DOGE, and LINK over a 90-day window in 2023-2025. The colored lines in the two sets of graphs illustrate how these correlations change over time, moving up or down as macro and crypto market conditions shift.

Practitioner-oriented platforms, such as TradingView, enable the construction of dashboards that depict multilateral relationships using sophisticated and visually compelling graphical representations.

Figure 7 above reveals the rapidly changing correlations between Bitcoin and other macro aggregates with correlations that fluctuate widely between -1 and +1. In the lower quadrant, most major cryptos (BTC, SOL, ETH, ADA, etc.) exhibit high inter-correlation, especially during risk-on or panic phases as presented with the “bell-shape” constructions. Periodic dips or shocks, potentially due to market events or macro shifts (e.g., regulatory news, federal fund rate expectations, liquidity crunches) can temporarily weaken correlations. The correlation table presents a more diversified interaction between so-called alt-coins and Bitcoin highlighting alternating “seasons” in favor of one or the other category.

3.3 Subjective Factors

“Perception is reality”. While no macro or finance aggregate definitively stand out as leading indicators, the common denominator of any valuation driver may be an overall subjective perception of markets formed by a vague perception of risk, wishful expectation of future returns or any “fear of missing out”. As noted by Beckmann et al. [BGW24] retail investors forming around 80% of crypto market participants are prone to stronger reactions to incoming public information and fragmented narratives because market sentiment and media attention largely motivate their investment decision. Valuations movements may be hence exacerbated due to psychological and cognitive biases that could result in suboptimal or irrational investment decisions. This could especially impact small crypto capitalization, alt- or meme coins exploited by orchestrated “pump and dump” schemes.

“Cut through the noise”. Irrational or “herd” behavior has a direct impact on the price formation mechanism, and the resulting “noise” can disconnect valuations from fundamentals if persistent and hence impact volatility [Sin21]. The challenge is then to “deflate” any signals from its noise component in order to extract valuable information with minimal distortion (decrease the noise to signal ratio). Such conditions undermine risk-premia-based arbitrage: rational, disciplined process-driven investors who rely on fundamental or technical analysis may encounter challenges executing strategies, as volatility and unpredictability increase. As noise traders (“buy the rumor, sell the news”) exacerbate market movements through rapid, sentiment-driven buying and selling, the resulting instability erodes price discovery mechanisms and could, in turn, weaken the market's attractiveness for arbitrage and long-term investment.

The role of social media. Cryptocurrency markets offer an ideal "playground" for quantifying the dynamic link between social media sentiment and price volatility, as they lack established valuation frameworks and regulatory oversight typically found in traditional finance, when it comes to information and research sources. In a context of behavioral biases that predominantly affect retail participants, the information flow in cryptocurrencies is dominated by social media platforms. Empirical research demonstrates that social sentiment, measured via platforms like Twitter, Telegram and Reddit, can exert a significant and immediate predictive effect in propagating price fluctuations, often outpacing the impact of conventional news sources and channels [KD20; La +21].

3.3.1 Managing Sentiment

Sentiment analysis is indispensable in managing cryptoassets due to properties such as:

- **Comprehensiveness:** Its holistic nature is the result of a whole array of heterogeneous variables and dimensions, perspective and influences (economic, financial, psychological or social)
- **Anticipative:** Sentiments are nurtured from anticipations or expectations related to future events should they materialize or not. In general, fear or any risk aversion is logically related to future and not past events. In that case sentiment analysis if translated into a single or set of figures is a forward looking indicator.
- **Synthetic score:** Subjective assessments and other contributing factors can be transformed into one single homogeneous figure such as the "Fear and Greed Index", "Social media engagement coefficients", "whale activity", "number of crypto-related tweets" or any other language related counts methodologies.

Several platforms[10] offer ready-made aggregates and analytics as sentiment measures as captured by social media activity surrounding certain topics can be predictive of their future performance or forthcoming investors attention.

3.3.2 Fear and Greed Indices

Given the increasingly close ties between finance and cryptoassets, it is both timely and relevant to compare the popular "Fear and Greed" indices used to gauge sentiment in traditional financial markets and digital asset sectors. These indices serve as widely recognized barometers of market sentiment, allowing investors to quantify collective emotion and better understand prevailing risk attitudes across both market segments.

Fear and Greed indices can refine buy-and-hold strategies by signaling contrarian opportunities, for instance when sentiment extremes are detected such as through RSI-type indicators. Investors can initiate trades that capitalize on exaggerated fear or greed levels such as presented by Let et al. [Let+23][11]. While such rule-based may perform better (risk-adjusted excess returns) than passive (buy-and-hold) strategies or pure opportunistic trading schemes, consistency and statistical significance may not be guaranteed due to the intrinsic weakness of "all-in" indicators. Reflecting overall market sentiment, the F&G index and similar indicators rely on publicly available data such as social media and surveys, which may

[10]For instance, Santiment aims to provide a concise overview of the most discussed topics currently in the crypto social media. For this purpose, the platform identifies every 60 minutes the top 10 words with the biggest spike in social media mentions compared to their average social volume over the previous 2 weeks. This signals an abnormally high interest in a previously uninspiring topic, making the list practical for discovering new and developing talking points in the crypto community. Using a similar methodology as Santiment, Augmento collects cryptocurrency related conversations from Twitter, Reddit & Bitcointalk and applies a classifier trained on crypto specific language before the data is analyzed according to 93 sentiments and topics.

[11]The tested strategy assumes that the contrarian investor buys the cryptocurrency when the Crypto Fear & Greed Index is characterized by low values, and sells it when the values are high, indicating that the market is due for a correction. As a robustness check, six levels of market entry and six levels of market exit are examined: Open a long position whenever the FGC index is sufficiently low (investor anxiety). In the study, six levels of entry are considered (FGC $\leq$ 20, FGC $\leq$ 21, FGC $\leq$ 22, FGC $\leq$ 23, FGC $\leq$ 24, FGC $\leq$ 25). Close the long position whenever the FGC index is sufficiently high (investor euphoria). In the study six levels of exit are considered (FGC $\geq$ 80, FGC $\geq$ 79, FGC $\geq$ 78, FGC $\geq$ 77, FGC $\geq$ 76, FGC $\geq$ 75). Passive rule: benchmark reference. The assumption is that a passive investor buys a cryptocurrency on a random day t and sells it on a day t + 365.

Figure 8: **Observations Fear & Greed**

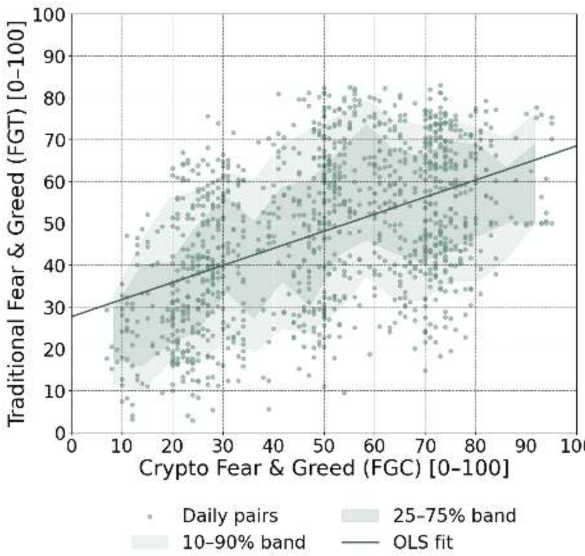

Figure 9: **Correlation history Fear & Greed**

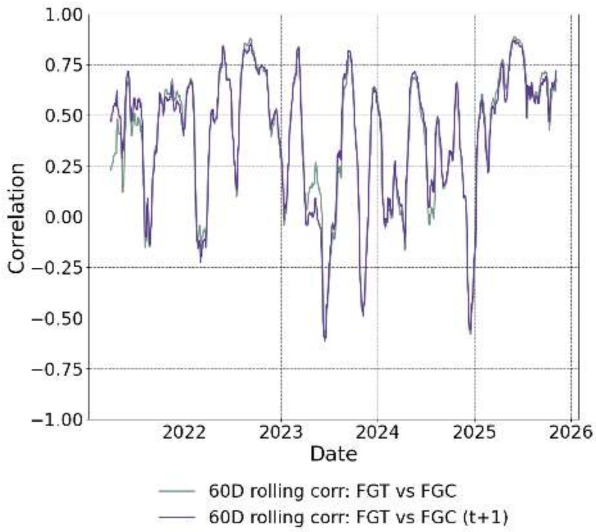

Note: Figure 8 shows realized values of the Fear & Greed Indices, the scattered points represent daily paired observations, while the shaded bands indicate percentile ranges (25–75% and 10–90%) and the solid line shows the OLS regression fit. Figure 9 shows the 60 day rolling correlations between the traditional Fear & Greed Index (FGT) to the Crypto Fear & Greed Index (FGC) and also to the FGC lagged by one day after.

omit important market information including idiosyncratic risks or token-specific news. They primarily capture short-term sentiment shifts rather than long-term trends, making them less useful for long-term investment decisions. Also, while they provide insight into market sentiment, they don't directly predict price movements. Both indices are reported on the conventional 0–100 scale, where lower values indicate "extreme fear" and higher values "extreme greed", and are observed at a daily frequency.

While trying to best capture investor sentiment, both indices are constructed using different components in terms of data sources and signal weighting, for the concrete construction see Appendix A.

Taken together, Figures 8 and 9 provide a compact view of how traditional and crypto Fear & Greed indices relate to one another. In Figure 8, the upward slope of the OLS fit indicates a positive correlation: as sentiment in crypto markets shifts from fear to greed, traditional market sentiment tends to move in the same direction, although not perfectly. The spread of points and the width of the bands signal substantial dispersion and volatility, particularly on the crypto side, consistent with the broader sentiment volatility discussed earlier. Figure 9 shows that, despite this noise, the rolling correlation between the two indices is predominantly positive, with only short-lived excursions towards zero or negative values. Overall, the charts visually reinforce that, despite structural and participant differences, investor sentiment in both domains is interconnected and frequently moves in tandem.

Figure 10 summarizes how the two sentiment indices co-move with global equity markets and how volatile their own dynamics are. The lower panel describes the second moment of "volatility" embedded in the construction of both indices, illustrating how unpredictable changes in market mood can become in response to news and macro events. Early 2022 shows both crypto and traditional finance experiencing elevated volatility, driven by the Russia–Ukraine war that began in February 2022 and the Federal Reserve's aggressive rate-hiking cycle from March 2022, with the traditional finance index (FGT) reacting particularly strongly. The crypto-specific spike from mid-2022 is attributable to the Terra/LUNA collapse in May 2022, while the subsequent collapse of FTX at the end of 2022 did not generate an additional, distinct volatility surge in the crypto Fear & Greed index (FGC).

The spike in April 2025 represents a synchronized shock affecting both asset classes: President Trump's announcement of sweeping tariffs on April 2, 2025 ("Liberation Day") triggered a broad risk-off move, and both FGC and FGT volatility spiked in tandem. This episode illustrates how macro shocks propagate simultaneously across traditional and crypto markets when the underlying driver is a global policy event.

In the upper panel of Figure 10, the FGT index generally exhibits a persistently positive 60-day correlation with world stock returns, mostly between 0 and 0.3, with brief periods closer to 0.5, especially

Figure 10: **Comparative correlations and volatility of the F&G Crypto and F&G Traditional**

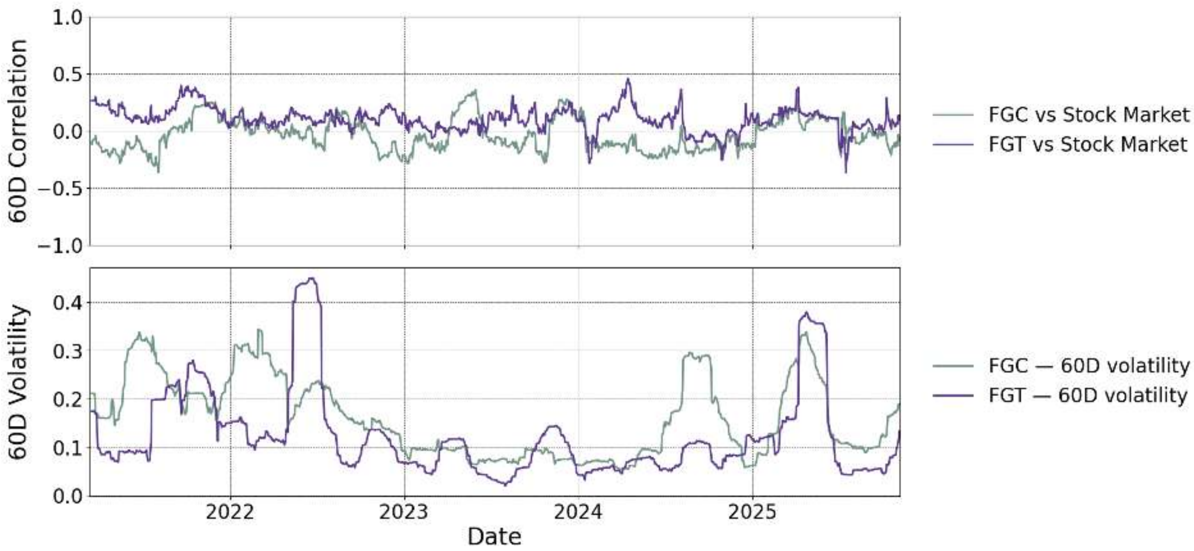

Note: FGC denotes the Crypto Fear & Greed Index and FGT the traditional Fear & Greed Index.The upper panel shows 60-day rolling correlations of both indices with global stock market returns; the lower panel shows 60-day trailing volatility of the two indices.The sample runs from 1 January 2021 to October 2025. All series are based on daily observations, and statistics are computed on overlapping 60-day windows.

around 2021. This indicates that when traditional finance sentiment is bullish (high greed) or bearish (high fear), global equities tend to move in the same direction, consistent with a synchronized sentiment–return relationship. In contrast, the FGC index shows a weaker and more volatile correlation with stock returns and more frequent dips into negative territory. This suggests that crypto market sentiment is less tightly anchored to global equity performance: there are intervals where crypto sentiment and stock returns are inversely related or largely disconnected, reflecting the more idiosyncratic nature of crypto-specific news and regime shifts.

Since crypto sentiment often diverges from traditional markets, the two sectors can sometimes offer diversification benefits. Both indices show periods of higher and lower correlation, hinting that during major global events or crises, there could be brief alignment between both crypto/traditional sentiment and stock performance. However, these do not persist and revert back to weak or no correlation.

3.4 Conclusions on Valuation Driving Factors

Macroeconomic factors influence cryptocurrency markets, but they do not fully account for price movements or investor behavior. The impact of economic indicators is inconsistent across different market regimes; their release may serve as catalysts for short-term rallies and shifting risk aversion, yet the underlying narrative can change with each episode, offering little continuity. The main transmission channel is risk aversion, which connects investor sentiment to short-term, thematic trading patterns. Emotional factors likely influence uncertainty more than the fundamental price component. When it comes to investor's behaviors, the underlying trader demographics and their financial literacy result in fundamentally different live sentiment profiles and volatility regimes. Crypto sentiment is more subject to sharp and rapid shifts, mirroring a retail-dominated and inherently volatile asset class, whereas traditional sentiment indices tend to reflect more consensus-driven and stabilized market conditions. As a result, the traditional Fear and Greed Index exhibits less erratic movement compared to its crypto counterparts, reflecting both market structure and participant sophistication. The reliance on real-time, community-driven sentiment amplifies herd behavior and overreaction, thereby magnifying speculative momentum from immediate reactions and short-term volatility in the absence of institutional grounding and stable fundamental anchors.

4 Portfolio Construction Engine

Having outlined the key valuation drivers in Section 3, we now turn to the mechanics of portfolio construction. Our framework builds on that conceptual foundation by integrating comprehensive data sources that connect those drivers to market behavior. Before delving into the specifics of the allocation methodology, we briefly review the key data sources and analytical dimensions that feed into our models. Readers interested in how the data influence portfolio outcomes may find the following subsection helpful, while those focused on allocation mechanics can proceed directly to the discussion of the Black–Litterman framework.

4.1 Data Integration and Analytical Dimensions

To build a robust crypto portfolio, we draw on a wide array of data sources that capture different dimensions of the digital asset ecosystem. This subsection summarizes how these data are integrated and highlights the key analytical metrics derived from them. While we focus on a representative subset of data streams, these sources are part of the broader set of factors discussed in Section 3.

Market capitalization and liquidity. We obtain a snapshot of the market universe from CoinGecko by ranking cryptocurrencies by market capitalization. On September 21, 2025, we recorded the top 500 tokens, representing more than 99.9% of the total market capitalization tracked by the aggregator. Our sample spans February 10, 2018 through September 21, 2025, covering both bull and bear phases. From CoinGecko's data we extract individual market capitalizations and use them to construct market-weighted indices and compute measures of relative size and liquidity concentration. The CoinGecko documentation notes that the platform aggregates data from more than 1,000 centralized exchanges and over 1,700 decentralized exchanges, providing comprehensive coverage of trading venues.

Price and volume data. For transaction-level price and volume data, we rely on Binance, the largest centralized exchange by trading volume. Binance's real-time OHLCV candles offer granular data across a broad set of trading pairs. We aggregate these candles into daily closing prices for the identified top cryptocurrencies, denominated in USDT. From these daily price series we calculate returns, variances and momentum measures that feed into the investor's view component and guide reallocation decisions.

Protocol analytics. Beyond price and capitalization, we gather protocol-level metrics from DefiLama. Key indicators include Total Value Locked (TVL), which measures the capital committed to decentralized finance protocols, and other metrics such as liquidity provisioning, staking rates and protocol governance activity. Together with price, volume and market capitalization, these variables form the raw inputs for our custom signals. We also derive additional indicators—such as market breadth, turnover ratios and cross-sector momentum—that help predict asset outperformance.

Integration with other factors. The data streams described above are integrated with factors discussed in Section Valuation Driving Factors. These include rating-based assessments, governance scores, on-chain activity metrics and technical indicators. The combined dataset underpins our multi-layered signals, which are then employed to compute custom views used for the portfolio allocation model. While the focus here is on a representative subset of data sources, our framework can accommodate additional datasets as needed.

Having established the data foundation, we now turn to the portfolio construction methodology itself. The next subsection describes how a Black–Litterman approach combines the market equilibrium with user-derived signals to produce robust, balanced portfolios.

4.2 Black–Litterman Framework

Building on the integrated data and analytical insights, our portfolio construction engine follows a Black–Litterman (BL) framework [BL92; HL99; Meu08; Fab+07]. It blends a baseline of expected returns with user-supplied information (signals or forecasts) to produce robust and stable portfolios. As sketched in Figure 11, the process starts from market (strategic) weights, infers implied equilibrium returns via reverse optimization, combines these with factor-based and discretionary views weighted by their confidence, and then solves a final constrained optimization to obtain implementable portfolio

Figure 11: **Schematic of the Black–Litterman inference process**.

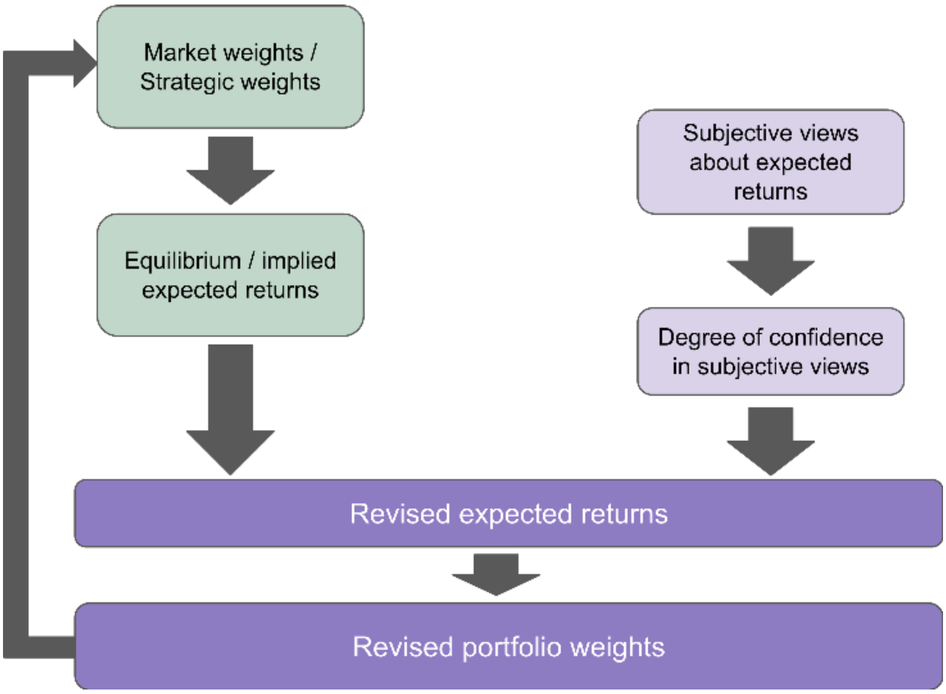

Note: Starting from market (strategic) weights, the model infers implied returns via reverse optimization, blends these with subjective views—derived from the major factors and their confidence—to obtain revised expected returns, then performs a final optimization to yield the revised portfolio weights.

weights. Unlike "naive" mean–variance optimization, which can be extremely sensitive to expected-return estimates and thus produce extreme weightings, the BL approach starts from a defensible equilibrium prior and then tilts the portfolio according to the investor's convictions, typically yielding smoother and more intuitive allocations [BL92; HL99; Fab+07].

Figure 12 contrasts reverse optimization with standard mean–variance optimization: in mean–variance, one supplies expected returns, correlations and a risk tolerance to obtain asset weights; reverse optimization takes current weights, correlations and a risk-aversion coefficient as inputs and infers the implied returns that would make those weights optimal. In our implementation, this implied-return vector forms the prior that is subsequently adjusted by the view layer before entering the final constrained optimization.

4.2.1 Input Options

The engine accepts three kinds of baseline inputs from the practitioner:

1. **Baseline market weights:** By default, the market capitalization vector is treated as a neutral "prior" portfolio. The engine applies reverse optimization (a CAPM-like calculation) to this weight vector to infer a set of expected returns that would make the market portfolio optimal[BL92; HL99]. These implied returns serve as the starting point for the BL update.

2. **Custom expected returns:** Practitioners who already have forecasts stemming from a factor model or a discretionary view may supply them directly. In this case, the BL update treats these forecasts as the baseline to be refined; the model will blend them with any additional views to produce the final posterior returns[Meu08; Fab+07].

3. **Custom target weights:** If a user has a heuristic allocation (for example, equal-weighted or risk parity), the engine can reverse-optimize those weights to derive implied returns and then proceed with the BL update. This allows the heuristic to be used as a prior while still incorporating new information.

Regardless of which baseline is chosen, practitioners can augment it with additional signals or forecasts

Figure 12: **Reverse optimization versus traditional mean–variance optimization**

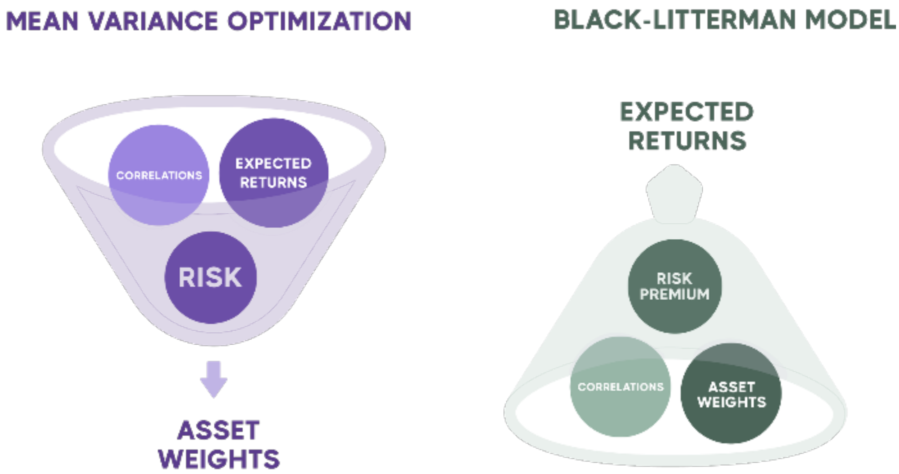

Note: The left funnel depicts mean–variance optimization (inputs: expected returns, correlations and risk tolerance; output: weights). The right funnel shows reverse optimization (inputs: current weights, correlations and risk tolerance; output: implied returns). Black–Litterman uses reverse optimization to derive the prior expected returns, then updates them with investor views.

in the form of views. Views may be absolute (e.g., "Ethereum will return 10% next year") or relative (e.g., "Solana will outperform Bitcoin by 2%"), and they can also be ordinal or rank-based when precise forecasts are hard to specify[Çel+21; Çel+25]. These investor views are distilled representations of the data described in Section 4.1 and the valuation drivers summarized earlier. Each view encapsulates multiple quantitative signals—such as relative momentum, market capitalization trends, trading volume and protocol metrics—into a single expectation about the relative performance of a subset of assets. By condensing the rich information from heterogeneous factors into a handful of directional opinions, the model reduces dimensionality and allows the user to express complex convictions in a tractable way. The choice of which signals to emphasize and how confidently to weight them represents a form of prior judgment.

When processed through the BL "machinery", these aggregated views shift the implied equilibrium returns toward the investor's synthesized outlook while preserving diversification benefits and maintaining consistency with the underlying correlations. Each view comes with a confidence level, allowing the model to weight it appropriately: strong views (high confidence) tilt the allocation more, while weak or uncertain views leave the portfolio close to the prior[BL92; Meu08]. For example, Figure 13 shows how introducing a single relative view, e.g., expecting Solana to outperform, tilts the baseline market-cap weight allocation: the favored asset's weight increases, while other weights decrease slightly to preserve diversification. This example illustrates how views are integrated in practice and underscores the importance of confidence levels in controlling the magnitude of the tilt.

Details on encoding views and their linear constraints are given in Appendix B. We will revisit the process workflow in Figure 11 when we discuss the update and optimization steps below.

4.2.2 From Baseline to Posterior Returns

After establishing the baseline and collecting views, the BL update produces a revised set of expected returns (the "posterior"). Conceptually, the posterior returns are a weighted average of the prior returns and the view-implied returns, with weights determined by the user's confidence and a scaling factor[BL92; Meu08]. This blending ensures that even small or uncertain views nudge the portfolio gently, while strong views can meaningfully tilt it. The update also propagates information through the covariance matrix, so that views on a few assets affect related assets via correlations[Fab+07]. This step corresponds to the second stage of the process workflow illustrated in Figure 13, where implied returns are updated with aggregated investor views.

Figure 13: **Example on how a single view tilts the portfolio weights**

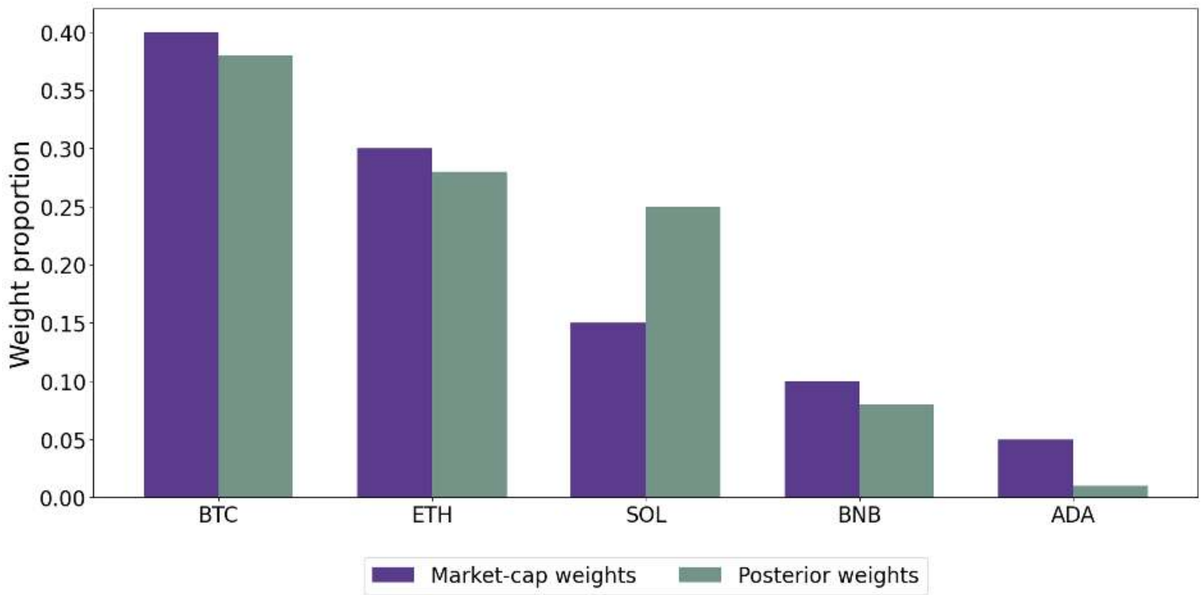

Note: Purple bars represent market capitalization weights; green bars show the posterior weights after assuming Solana will outperform. The view increases the weight of Solana and reduces the weights of other assets. Weights remain diversified because confidence levels control the magnitude of the tilt.

4.2.3 Optimization and Constraints

The final stage of the process workflow (Figure 11) solves a constrained mean–variance problem, maximizing expected return for a given level of variance subject to:

- **Long-only positions** (no short selling).
- **Full investment**, ensuring the weights sum to one.
- **Per-asset floors and caps** to avoid extreme allocations and ensure diversification.
- **Additional group or sector limits** if required by mandate.

To account for estimation error, we add a robustness penalty that discourages positions relying heavily on uncertain return forecasts, reducing concentration and improving out-of-sample stability[Fab+07; BBC11]. The resulting portfolio stays close to the baseline (e.g., market-cap weights) but tilts toward the investor's views in proportion to confidence. The same convex formulation can also include additional linear or quadratic constraints, such as risk-budget or sector-contribution limits; full details are given in Appendix B.

4.3 Results

We next present empirical results and performance charts, comparing the BL strategy with the following benchmark portfolios:

- **Black–Litterman** and **Black–Litterman (with Fee)** are our proposed strategy, reported before and after incorporating transaction-cost adjustments.
- **Market Cap Top 10** is the market-cap-weighted portfolio of the ten largest assets.
- **Market Cap All** holds all investable assets in proportion to their market capitalization and serves as the broad market proxy.
- **Market Cap Selected** is the market-cap-weighted portfolio restricted to the subset of assets selected by our strategy (Section 3).
- **Equal Weight Top 10**, **Equal Weight All**, and **Equal Weight Selected** assign equal weights over the same universes as the corresponding market-cap portfolios; **Equal Weight All (with**

Fee) additionally applies the transaction-cost model used in our backtests.

- **BTC** and **ETH** are single-asset buy-and-hold allocations to Bitcoin and Ether, respectively.
- **CCI30** is an external diversified crypto index used as an additional benchmark.

Figure 14: **Cumulative returns of BL strategy and benchmarks**

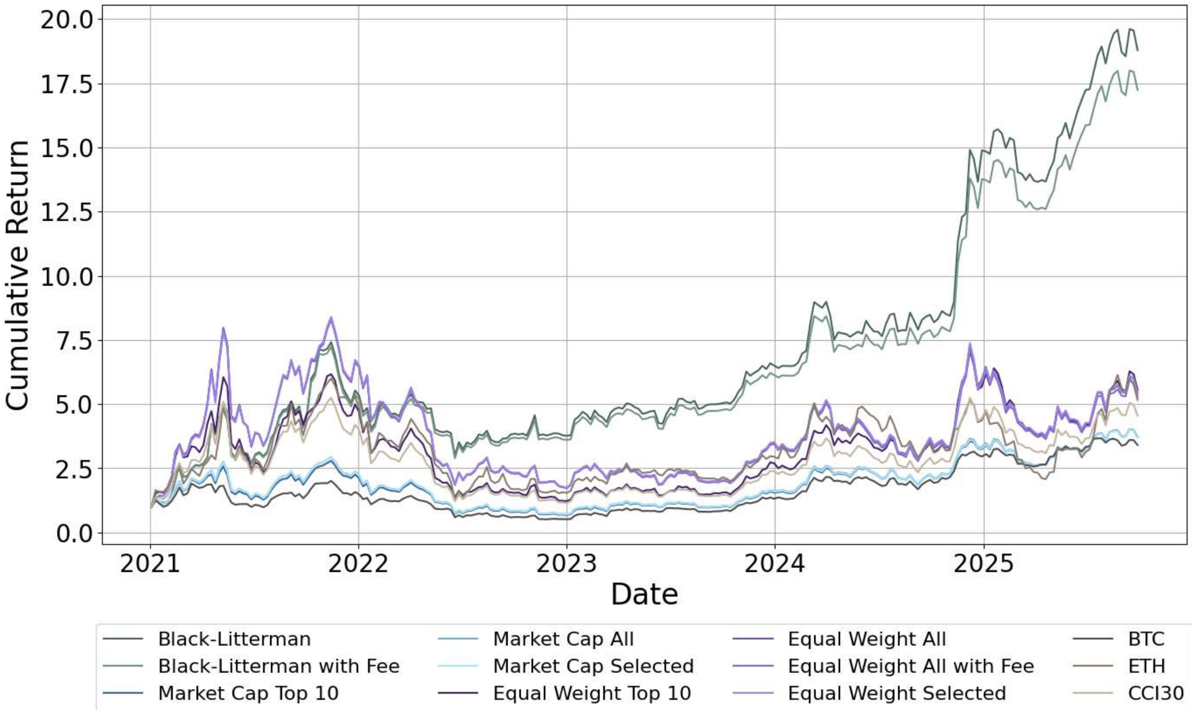

Note: BL and benchmarks rebalanced weekly from Jan 2021 – Sep 2025

4.3.1 Overall and Downside Risk-Adjusted Performance

Figure 14 reports cumulative wealth for the BL strategy and benchmark portfolios from Jan 2021 to Sep 2025 under weekly rebalancing. Starting from a value of 1, the baseline BL portfolio reaches nearly 19× over the sample, while the fee-adjusted variant ends around 17×. In comparison, market-cap portfolios (Top 10, All, Selected) finish near 3.7×, and equal-weighted variants around 5.4×. The fee-adjusted BL path closely tracks the frictionless BL series but at a slightly lower level, reflecting the impact of realistic trading costs. Detailed return and risk measures corresponding to these paths are summarized in Table 3.

To track performance relative to downside risk over time, we report a 12-month rolling Sortino ratio (annualized) for the BL strategy in Figure 15. Crypto return distributions are well-documented to be non-normal, with heavy tails, volatility clustering, and asymmetric tail risk concentrated on losses. Empirical studies find fat-tailed downside events in major coins [GK18; Kat17] and show that bad (downside) volatility shocks dominate risk spillovers across the market, whereas good (upside) volatility is less informative for risk transmission [Gha21]. In line with broader evidence that semivariance isolates economically meaningful downside variation [And+12], we therefore focus on a downside-based measure (Sortino) rather than total-volatility measures (Sharpe) for these rolling diagnostics.

Empirically, the BL strategy's rolling Sortino ratio remains positive over the full sample and spends most of the time at elevated levels (typically above 2), with a pronounced dip towards lower values around the end of 2022 during the major market drawdown. Subsequent windows recover quickly to levels comparable to, or higher than, those observed in the early part of the sample. Figure 15 thus complements the cumulative-wealth plot by showing that high long-run returns are accompanied by consistently favorable downside risk-adjusted performance, punctuated by a short period of deterioration during extreme market conditions.

Figure 15: **Rolling 12-month Sortino ratio of BL strategy (annualized)**

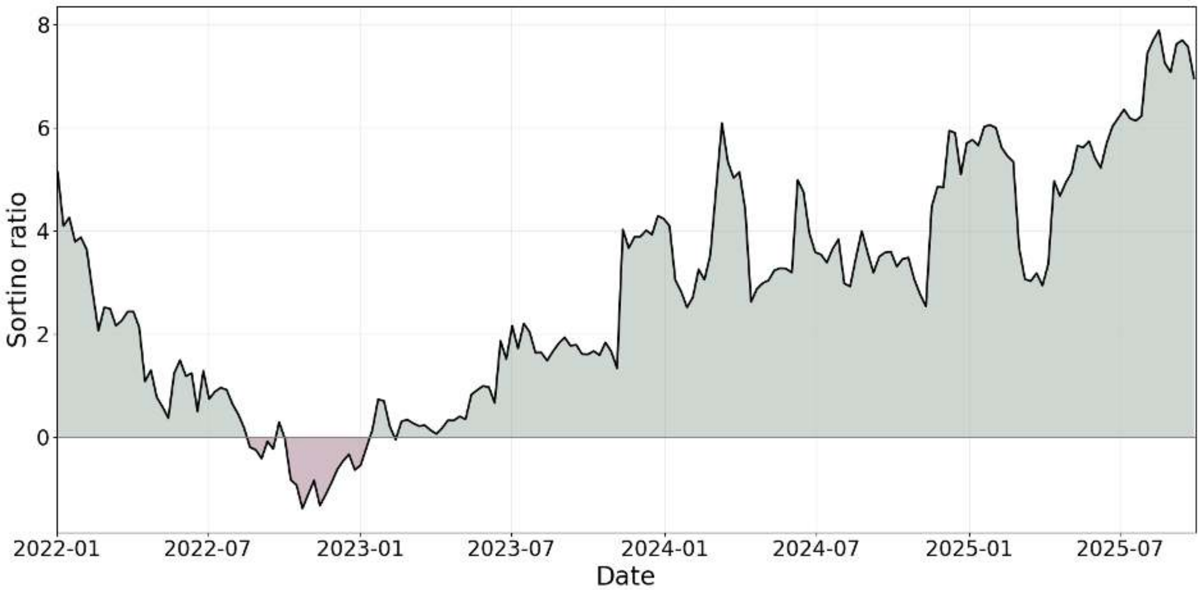

Note: Annualized 12-month rolling Sortino ratio of the BL strategy over Jan 2021 – Sep 2025, computed from weekly returns using downside deviation based on negative observations only.

4.3.2 Relative and Regime-Specific Performance

Figure 16 summarizes the strategy's performance versus the CCI30 benchmark at a monthly and cumulative horizon. The bars show monthly excess returns of the BL strategy relative to CCI30, while the line tracks cumulative relative wealth rebased to 1 at the start. Individual months alternate between positive and negative excess returns, including several periods with notable short-term underperformance. Over the full sample, however, the cumulative relative-wealth curve trends upward and ends well above 1, indicating that positive excess months dominate sufficiently to generate a persistent performance differential.

Figure 16: **Monthly excess returns and cumulative relative wealth vs. CCI30**

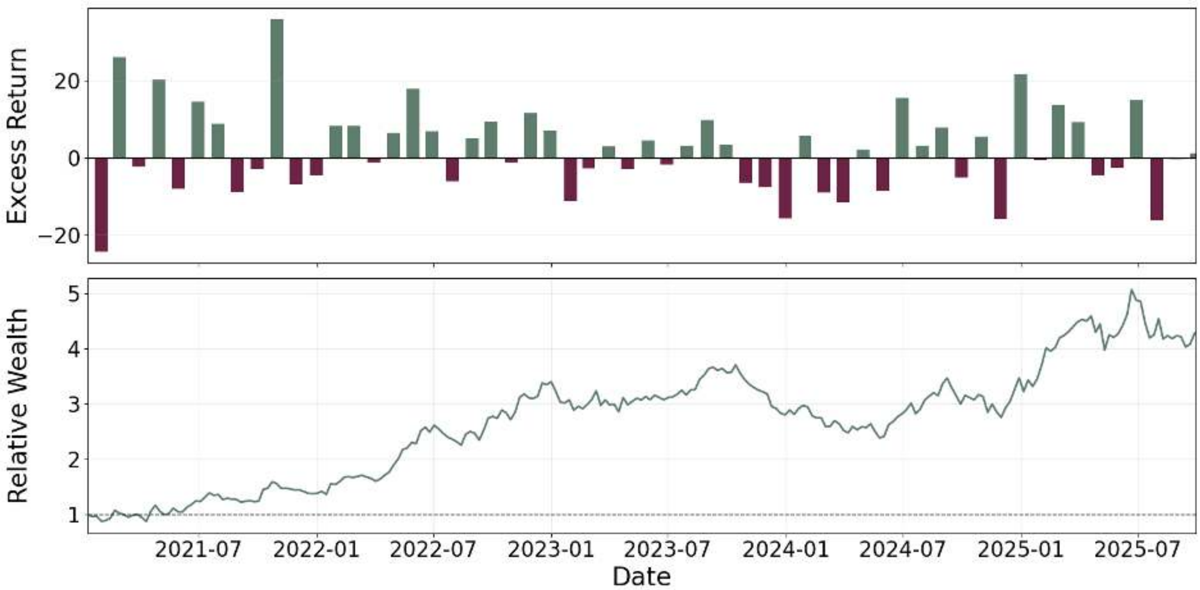

Note: Monthly excess returns (bars) and cumulative relative wealth (line) of the strategy vs. CCI30. Bars above zero indicate months the strategy outperformed CCI30; the line tracks the compounded outperformance over time (rebased to 1 at the start).

Figure 17: **Cumulative returns by calendar year for BL strategy and benchmarks**

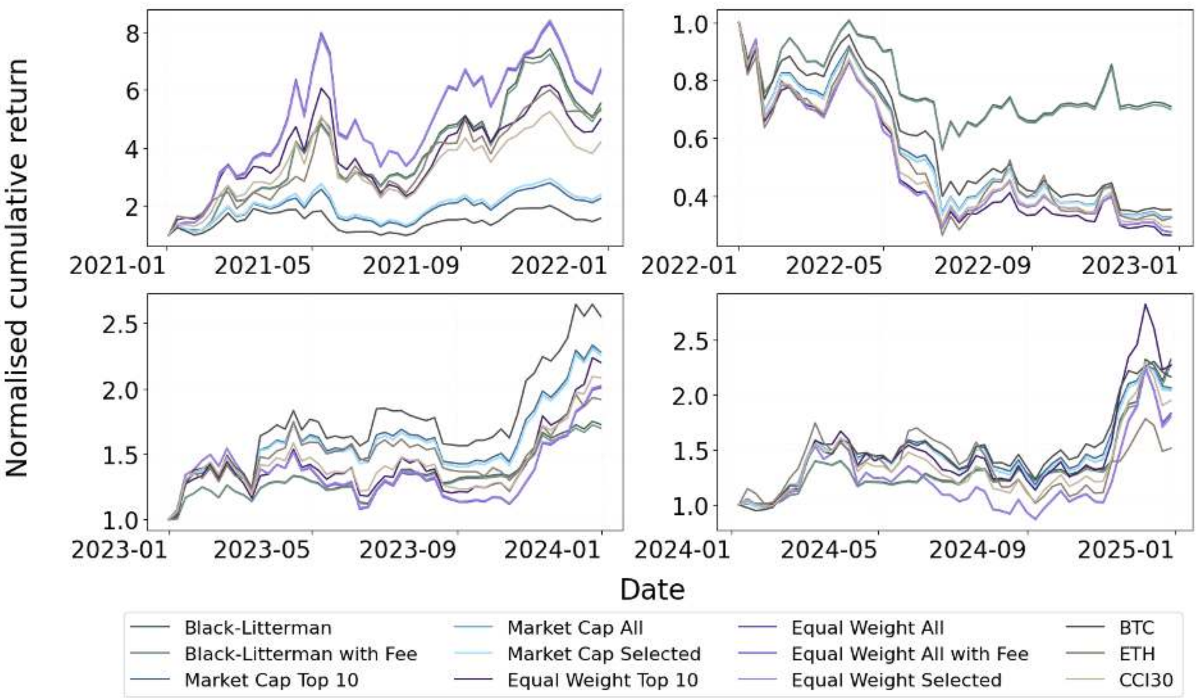

Note: Cumulative returns by calendar year (2021–2024) for BL and benchmark portfolios, with weekly rebalancing.

To examine behavior across broader market conditions, Figure 17 decomposes cumulative returns by calendar year, grouping the sample into four regimes: 2021 (bull), 2022 (bear), 2023 (recovery), and 2024 (sideways to trending). In 2021's strong up-market, the BL strategy is among the portfolios with the highest return in the comparison set. During the 2022 drawdown, all strategies experience substantial losses; the BL series records a sizable decline but remains within the upper part of the range spanned by the benchmarks. In the subsequent recovery and mixed conditions of 2023–2024, BL continues to rank near the top of the cumulative-return distribution. Within each year, the dispersion across benchmarks is substantial, highlighting that the ranking of simple market-cap and equal-weight strategies is sensitive to the prevailing regime. By contrast, the BL path moves more smoothly through these transitions, reflecting the combination of a diversified baseline and systematically updated views. Figure 17 documents how the strategy behaves across distinct market regimes, making explicit that both favorable and challenging phases are reflected in the realized path.

4.3.3 Summary Statistics

Table 3 consolidates absolute and risk-adjusted metrics for all strategies. It reports annualized arithmetic returns, total volatility, downside volatility, Sharpe and Sortino ratios (assuming a zero risk-free rate), beta relative to the broad market-cap universe ("Market Cap All"), and maximum peak-to-trough drawdown over the sample.

Within this set of strategies, the BL portfolio attains the highest annualized return and the highest Sharpe and Sortino ratios. The fee-adjusted BL variant reports slightly lower return but remains near the top across risk-adjusted measures. Benchmark portfolios exhibit lower average returns and, in several cases, higher volatility and deeper maximum drawdowns. Single-asset exposures (BTC, ETH) show comparatively high volatility and drawdowns relative to their Sharpe and Sortino ratios, providing context for the multi-asset allocations. The "Average" and "Median" rows further summarize the cross-strategy distribution: both BL variants sit well above these aggregate reference levels in terms of return and risk-adjusted metrics, while lying below them in terms of volatility, downside volatility, beta, and maximum drawdown.

Table 3: **Performance summary of BL strategy and benchmarks**

Strategy	Return	Volatility	Down Vol.	Sharpe	Sortino	Beta	Max DD
Black–Litterman	**78.91 %**	59.94 %	**28.75 %**	**1.3165**	**2.7450**	**0.7938**	**−59.63 %**
Black–Litterman (with Fee)	**77.04 %**	**59.88 %**	**28.84 %**	**1.2866**	**2.6710**	**0.7932**	**−60.08 %**
Market Cap Top 10	**46.76 %**	61.57 %	34.75 %	0.7594	1.3456	0.9883	−75.54 %
Market Cap All	47.11 %	62.24 %	35.12 %	0.7570	1.3414	1.0000	−75.87 %
Market Cap Selected	46.93 %	62.01 %	35.02 %	**0.7568**	**1.3399**	0.9962	−75.83 %
Equal Weight Top 10	64.70 %	76.36 %	40.05 %	0.8473	1.6153	**1.1399**	**−80.17 %**
Equal Weight All	64.88 %	**77.09 %**	40.58 %	0.8416	1.5987	1.1301	−79.20 %
Equal Weight All (with Fee)	64.30 %	77.07 %	**40.60 %**	0.8344	1.5838	1.1298	**−79.29 %**
Equal Weight Selected	64.76 %	76.69 %	40.56 %	0.8444	1.5997	1.1292	−79.28 %
BTC	**43.06 %**	**58.85 %**	32.63 %	**0.7317**	**1.3197**	0.9093	−74.55 %
ETH	67.98 %	**83.23 %**	**42.58 %**	0.8168	1.5965	**1.1837**	−78.67 %
CCI30	55.98 %	69.47 %	37.82 %	0.8059	1.4803	1.0798	−78.19 %
Average	60.20 %	68.70 %	36.44 %	0.8832	1.6864	1.0228	−74.69 %
Median	64.50 %	65.85 %	36.47 %	0.8256	1.5901	1.0399	−77.03 %

Notes: Measured Jan 2021–Sep 2025, weekly rebalancing. Returns, Volatility, Downside Volatility, and Max DD are shown in percent (annualized for the first three; Max DD is peak-to-trough over the sample). Sharpe/Sortino assume a zero risk-free rate. Beta is computed vs. "Market Cap All". Best two values are highlighted in green and worst two in red per column; "best" means higher for Return/Sharpe/Sortino and closer to zero for Vol/Down Vol/Beta/Max DD. The final two rows report the simple cross-strategy mean and median of each column; they are descriptive summaries and do not correspond to investable portfolios.

4.4 Interpretation

Taken together, Figures 14–17 and Table 3 yield three main conclusions.

First, the BL engine delivers higher cumulative growth than crypto benchmarks while taking less systematic risk. BL and its fee-adjusted variant sit at the upper end of the return spectrum, yet with beta below one and volatility comparable to or below broad market-cap and equal-weight portfolios. This suggests that the strategy does not simply scale up market exposure but adds an active return component on top of a diversified baseline.

Second, the path of returns reflects systematic management of downside risk rather than purely higher average returns. The rolling Sortino ratios remain elevated relative to benchmarks, indicating more return per unit of downside volatility. Monthly excess returns alternate between gains and losses, but the associated relative-wealth curve trends steadily upward, consistent with positive excess months accumulating over time. Calendar-year panels show both strong and weak years, including a pronounced drawdown in 2022, yet the BL portfolio tends to experience shallower peak-to-trough losses than many references. These elements jointly point to a profile in which downside episodes are mitigated more effectively while upside participation is largely preserved.

Third, the results appear robust to realistic implementation costs and to alternative sample splits. Introducing conservative trading fees compresses performance only moderately and leaves the ranking across Sharpe, Sortino, beta, and drawdown broadly unchanged. This suggests that the performance is not driven by a few unusually favorable periods or by frictionless assumptions. Instead, the evidence supports the core premise of the paper: a multi-layer, multi-factor Black–Litterman framework can aggregate heterogeneous valuation drivers into diversified, implementable allocations that improve both absolute and downside-risk-adjusted performance relative to passive crypto exposures.

5 Rebalancing Execution

The new target allocation weights generated by the portfolio construction engine reflect the general individual risk awareness, personal views on the uncertainty of the crypto market, and custom assessments in terms of expected returns for individual positions. Realizing a reallocation from the present to the new allocation comprises (i) the determination of the reallocation amounts, (ii) the cost-optimal calculation of orders and sequence of orders for those amounts, and (iii) the execution of the sequential orders.

5.1 Cost of Reallocation

In traditional finance, the sale of a position, if held in a taxable account, is usually a taxable event (financial transaction tax (FTT), also known as the securities transaction tax (STT)). Eliminating the cost of time and labor through automation, taxes on capital gains are the dominant cost of a reallocation (or rebalancing) event among other transactions fees. Zhang et al. [Zha+22] therefore suggest a careful selection of the rebalancing frequency. In crypto markets, on the other hand, those tax costs can be evaded in several jurisdictions[12]. In terms of trading, crypto is often embedded in countries' laws as a property or investment asset whose capital gains are taxed in the event of selling against cash or paying for services or goods. However, trading cryptocurrency against cryptocurrency directly means staying invested in crypto; no taxable event is triggered.

With the omission of the tax cost, the transaction cost is the relevant factor if a reallocation is economically reasonable [PS88; Lel99; HH12]. Those transaction costs can be minimized using a dynamic programming approach [Bel57; BD62; Ber00] which discovers the set of orders that satisfies the minimum-cost objective.

5.2 Graph Topology for the Reallocation of Cryptoassets

Stablecoins like USDT serving as a substitute for fiat money allow the reallocation of crypto portfolios in a similar way to how common stock portfolios are rebalanced. Overweight cryptocurrency positions w.r.t. the new allocation target are sold against a stablecoin, and underweight positions are increased to reach the target allocation. However in the crypto space, we are not necessarily limited to stablecoins as intermediary currency. Other cryptocurrencies are often quoted against the top cryptocurrencies, offering a broader range of options to sell or buy one position against another. Due to this fact, a purely intermediary currency which serves as a transitional node may not even be necessary. Those options expand further with exchanges offering better fees, when using the platform's native cryptocurrency for trading. Hence, the fundamental optimization task of reallocating or rebalancing a crypto portfolio is to find the cost-efficient sequence of orders to arrive at the target allocation.

Figure 18a shows typical rebalancing setups in the traditional finance and crypto domain. The weights of Assets 1 and 2 are decreased whereas the weights of Assets 3 and 4 should increase to achieve the target allocation. While all trade paths are routed via a common intermediary asset in the traditional case, a fully connected graph network without intermediary asset as in Figure 18b is theoretically possible in the case of cryptocurrency trading. At first glance, such a graph seems counterintuitive because it allows, e.g., selling Asset 1 for Asset 2, so basically buying an asset that is ultimately also meant to be sold to achieve the target allocation. But trading fees for selling Asset 1 for Asset 2 followed by selling the now increased position of Asset 2 for Asset 3 or 4 may offer lower fees than directly selling Asset 1 for Asset 3 or 4. Similar reasoning applies to trades on the receiving buy side.

Figure 18: **Theoretical trade paths of assets**

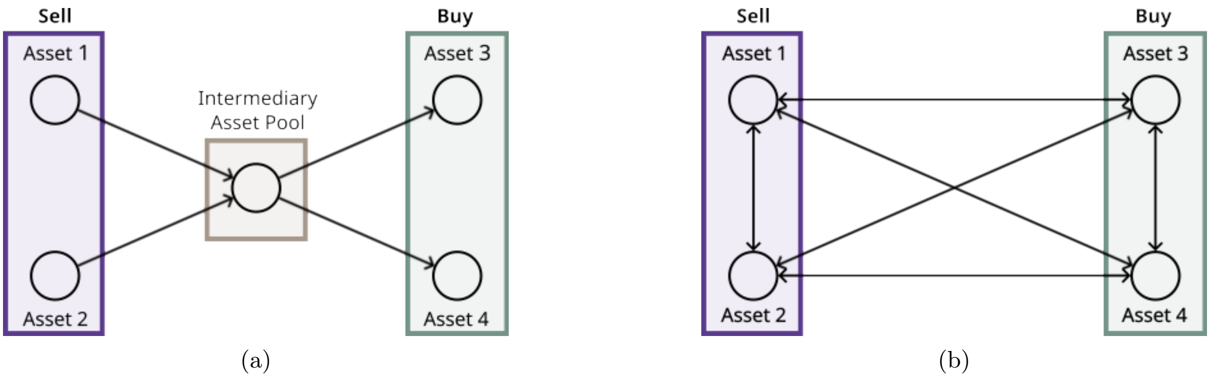

Note: (a) common stocks in traditional finance and (b) cryptocurrencies in the crypto markets.

[12] Simmons & Simmons provides a survey on crypto taxation in selected countries [Sim25]

5.3 Practical Considerations

In practice, the optimal sequence of orders is also affected by the price volatility of the trading pair [BL24] and the liquidity of the market [Bra+21; FM24]. Higher price volatility leads to greater slippage, that is, the price change between the order calculation and the actual execution of the recommended orders. An illiquid trading pair with limited demand and supply, on the other hand, may impair the fast sequential execution of orders in the sequence. The latter problem particularly applies to large position changes in a portfolio. However, the danger of cascading effects from sequentially executed orders was negligible in our reallocation tests, because trade networks with higher connectivity between assets shorten the average path length [Har69]. Reallocation order simulations and live execution of orders through our cost-optimal solution never reached a path length greater than 4. This finding also held for setups in which the path generator searched not only for trade paths within a single exchange, but across multiple accounts on different crypto exchanges[13].

5.4 Results

To determine the expected transaction cost savings of our cost-optimal approach using a complex network like in Figure 18b compared to the naive rebalancing approach through a common asset, we simulated the reallocation of portfolios on Binance exchange. From a random initial allocation, portfolios with a varying size and composition of assets [14] were reallocated to a new random target allocation. Both approaches used identical bid and ask prices and the same initial and target allocations. In the naive case, stablecoin USDC was used as a common asset, whereas the cost-optimal approach used all available routes between the assets in the portfolio composition and additionally routes to Bitcoin, BNB, Ethereum and USDC no matter if those assets are part of the portfolio or not.

Figure 19: **Cost savings of cost-optimal vs. naive approach**

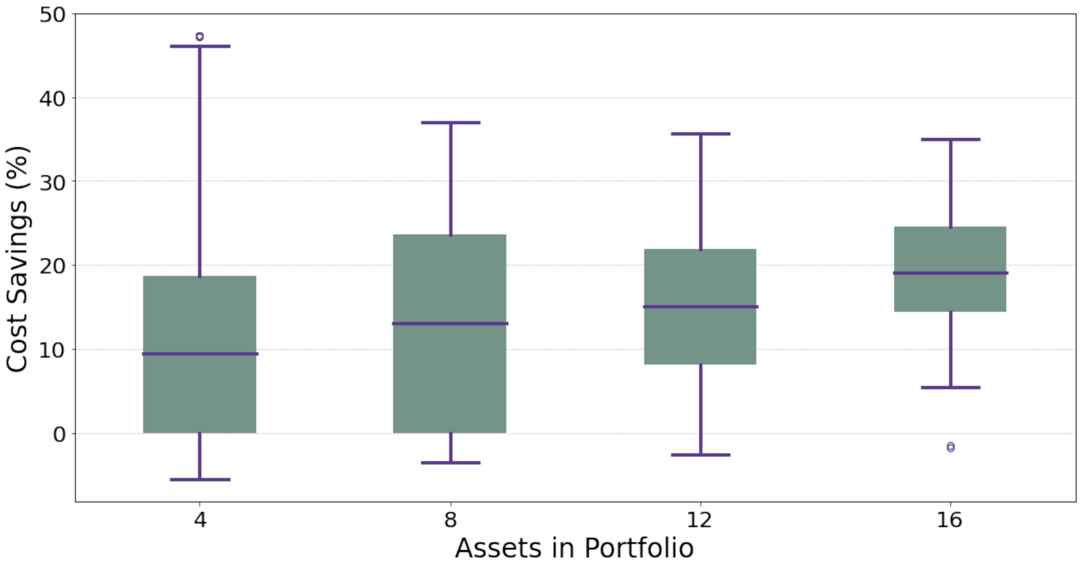

Note: Box plots for 100 simulations for each portfolio size with indicated mean. Whiskers are 1.5 times the interquartile range as defined by [Tuk77].

Figure 19 shows the superior performance of our cost-optimal approach. The mean cost savings for a 4-asset portfolio are 9.43% compared to naive reallocation and transaction cost savings increase for larger portfolios. Whereas the cost-optimal approach fails in a few cases, our solution outperforms the naive approach in 75 out of 100 cases even for the 4-asset portfolio. For the 16-asset portfolio the outperformance ratio increases to 99% and mean cost savings increase to 19,05%.

[13] The used exchanges were Binance, Bitfinex, Kraken, and Poloniex.

[14] The investment universe comprised 24 assets: Aave, Algorand, Avalanche, Bitcoin, Bitcoin Cash, BNB, Cardano, Chainlink, Cosmos, Dogecoin, Ethereum, Ethereum Classic, Hedera, Litecoin, NEAR Protocol, Polygon, Polkadot, Pepe, Shiba Inu, Solana, Stellar, Tron, Toncoin and XRP.

6 Discussion

Investment Process and Performance: "Two Sides of the Same Coin"

A replicable and clear investment process distinguishes disciplined investing from opportunistic speculation. This distinction becomes tangible through widespread systematic performance attribution framework [BF85] which decomposes portfolio returns into discrete, measurable components such as: allocation effects (how much capital to deploy across asset classes), selection effects (the choice of specific securities), and interaction effects (the combined impact of these decisions). By segregating these sources of return, an investment process enables managers and stakeholders to precisely identify whether outperformance stems from any of these aforementioned dimensions in addition to overall market movement and have the possibility to adjust accordingly as depicted in Figure 2.

Our approach demonstrates full compatibility with segregated performance analytical approaches, enabling transparent decomposition of return sources across contributing factors which differ in nature, velocity and magnitude. The imperative for a replicable investment process becomes even more pronounced in the cryptocurrency investment realm, characterized by recurring heightened volatility and sharp transient market dynamics. Rather than reacting to market noise or pursuing short-term trend-based strategies, a replicable process ensures that investment decisions are accountable and risk-managed without making any structural concessions to market phases (bull run, meltdown or "crypto winters").

The "Jazz" Trio: Macro, Crypto, Sentiments

Predicting the degree and duration of macro effects on crypto remains complex. No single indicator or data point reliably drives valuations over time. While crypto markets increasingly react to macroeconomic development via institutional investors, establishing themselves as prominent actors, the relationship remains erratic and asymmetric. Crypto markets tend to react to macro drivers only during specific periods of heightened attention, resulting in an inconsistent and asymmetric relationship.

Emotional factors cannot be ignored, as they could lead to measurement errors or biases. Refining feature extraction for sentiment is crucial, given the low signal-to-noise ratio of emotional indicators. Advanced methodologies like machine learning can be highly effective in capturing fragmented and decentralized market sentiment. Providing additional investment signals, AI could be further trained on platforms like Telegram, YouTube or any dedicated forum, where a majority of retail investors form their opinions.

As a result, the effectiveness of these sophisticated tools should be approached with caution, given the prevalence of informal information sources dominated by retail investors. In this context normative and "semi-static" measures such as ratings could provide stability and references within an investment process, namely when it comes to ex-ante exclusion and exclusion criteria based on tangible facts and figures. However behavioral finance research demonstrates that investors often use any kind of ratings mechanically without fully understanding the underlying methodologies or its limitations [BT03]. In crypto markets characterized social media influence, and "fear of missing out" dynamics, ratings may simply provide post-hoc rationalization for emotionally driven investment decisions rather than disciplined normative guidance.

Portfolio Construction and Execution: a (Re)balancing Act

The results in Section 4 indicate that a multi-layer Black–Litterman (BL) engine can translate heterogeneous crypto-specific information into allocations that are both intuitive and quantitatively robust. Starting from market-cap weights as an equilibrium prior and tilting with factor-based views, the BL portfolios deliver higher long-run growth than market-cap, equal-weight and single-asset benchmarks, while maintaining comparable or lower volatility, downside volatility, beta and drawdowns. The rolling Sortino diagnostics and regime-wise breakdowns suggest that this is not driven by a single favorable episode, but by a persistent improvement in downside risk-adjusted performance. This is consistent with the original rationale of BL in traditional asset management [BL92; HL99; Meu08; Fab+07], but here implemented with crypto-specific signals and constraints.

Section 5 shows that these allocation properties are meaningful only when combined with an explicit execution layer. By modeling rebalancing trades using historical and real-time order-book data from centralized venues, the framework incorporates fees, available depth and partial fills when estimating transaction costs. Comparing hub-and-spoke routing with graph-based routing over the full set of trading pairs illustrates that exploiting routing optionality can materially reduce simulated costs while keeping effective path lengths short. This contrasts with much of the empirical crypto-asset literature, which often abstracts from detailed transaction costs or venue structure and focuses primarily on allocation rules or return properties [BM19; GK18; Kat17].

Several limitations qualify these findings. The BL implementation relies on a particular factor specification and parametrization, and currently uses only a subset of the valuation drivers introduced in the framework; different factor mixes, confidence structures or risk budgets may lead to different trade-offs between return, risk and turnover. On the execution side, the analysis focuses on centralized exchanges and leaves the integration of on-chain execution (e.g., automated market makers or intent-based routing) for future work. Nonetheless, taken together, the portfolio engine and execution module illustrate how allocation and rebalancing can form an integrated component of a full investment process suitable for institutional adoption.

7 Conclusion

If the world of cryptoassets feels overwhelming with its complexity, unpredictability, and chorus of discordant voices, the investment portfolio is the place where everything comes together in harmony. With this in mind, we presented a comprehensive, institutional-grade framework for cryptoasset portfolio management, integrating the full investment life cycle from universe definition and factor construction to portfolio optimization and cost-optimal execution. Our methodology and application combine heterogeneous and fragmented data sources into a unified multi-factor architecture. Our Black Litterman-based construction engine provides a structured mechanism to translate complex signals and subjective convictions into stable asset allocations, while explicitly managing estimation error, diversification, and risk contributions.

The empirical evidence demonstrates that this approach can achieve strong absolute and risk-adjusted performance across diverse market regimes, outperforming market-cap and equal-weight benchmarks as well as the CCI30 index or a single investment in Bitcoin, even when realistic trading costs are incorporated. Rebalancing execution plays a decisive role in ensuring that these modeled allocations translate consistently into realized portfolio outcomes. By optimizing trading paths across a multi-asset graph and incorporating venue-specific fee structures, the execution layer enables reallocation to be implemented efficiently under real-world market conditions with scalable extensions to accommodate mandatory investment guidelines or self-assigned risk limits.

To take an analogy with a car, the motor and its related control systems are all interconnected: many parts like the clutch and transmission, brakes and accelerator act independently or in alternating cycles rather than working together all the time. The same applies in an investment engine where all valuation drivers can take the center stage, (de)correlate with each other, become completely muted for some time or just disappear. For the same reasons our framework is intentionally modular: each component such as the investment universe definition or screening, valuation drivers, signal aggregation, optimization, and execution can operate on a standalone basis (or sub-routine) while remaining an integral part of an organized workflow.

In addition to supporting governance, scalability, and transparent performance attribution, the framework accommodates both “rule-based” and more “discretionary” investment approaches. Among the collection of factors, processed signals or subjective components, the challenge lies in identifying which factors are likely to become dominant (or considered as such) and positioning the portfolio accordingly. This could be addressed by an investment committee which might favor a holistic or research-based approach or a more algorithmic-driven process powered by artificial intelligence orchestrating all dimensions (or a mix of both). In that context, our results highlight that crypto markets remain heavily influenced by subjective sentiment, regime shifts, and non-linear feedback loops, underscoring the necessity of integrating both quantitative signals and qualitative judgment.

Ultimately, the integrated investment process and portfolio application presented here offer a structured framework aligned with established best practices and provide a foundation for developing the next generation of digital-asset products.

References

[Aba+20] Emmanuel Joel Aikins Abakah, Luis Alberiko Gil-Alana, Godfrey Madigu, and Fatima Romero-Rojo. "Volatility Persistence in Cryptocurrency Markets under Structural Breaks". In: *International Review of Economics & Finance* 69 (Sept. 2020), pp. 680–691. DOI: 10.1016/j.iref.2020.06.035. URL: https://doi.org/10.1016/j.iref.2020.06.035.

[AMŠ21] Zdravka Aljinović, Branka Marasović, and Tea Šestanović. "Cryptocurrency Portfolio Selection—A Multicriteria Approach". In: *Mathematics* 9.14 (July 2021), p. 1677. ISSN: 2227-7390. DOI: 10.3390/math9141677.

[Alt68] Edward I. Altman. "Financial Ratios, Discriminant Analysis and the Prediction of Corporate Bankruptcy". In: *The Journal of Finance* 23.4 (Sept. 1968), pp. 589–609. DOI: 10.1111/j.1540-6261.1968.tb00843.x.

[And+12] Torben G. Andersen, Tim Bollerslev, Xin Huang, and Jose Vega. "Realized Semivariance". In: *The Annals of Applied Statistics* 6.4 (2012). URL: https://projecteuclid.org/journals/annals-of-applied-statistics/volume-6/issue-4/Realized-Semivariance/10.1214/12-AOAS561.full.

[BT03] Nicholas Barberis and Richard H. Thaler. "A Survey of Behavioral Finance". In: *Handbook of the Economics of Finance*. Ed. by George M. Constantinides, Milton Harris, and René M. Stulz. Vol. 1B. Elsevier, 2003, pp. 1053–1128.

[BK20] Silvia Bartolucci and Andrei Kirilenko. "A model of the optimal selection of crypto assets". In: *Royal Society Open Science* 7.8 (Aug. 2020). Open access under CC BY licence, p. 191863. ISSN: 2054-5703. DOI: 10.1098/rsos.191863.

[BGW24] J. Beckmann, T. Geldner, and J. Wüstenfeld. "The relevance of media sentiment for small and large scale Bitcoin investors". In: *Journal of International Financial Markets, Institutions and Money* 92 (2024), p. 101963.

[Bel57] Richard Bellman. *Dynamic Programming*. Princeton, NJ: Princeton University Press, 1957.

[BD62] Richard Bellman and Stuart E. Dreyfus. *Applied Dynamic Programming*. Princeton, NJ: Princeton University Press, 1962.

[BR23] Gianluca Benigno and Carlo Rosa. *The Bitcoin–Macro Disconnect*. Staff Report 1052. JEL: F3, F4, G1. Federal Reserve Bank of New York, Feb. 2023. URL: https://www.newyorkfed.org/research/staff_reports/sr1052.

[Ber00] Dimitri P. Bertsekas. *Dynamic Programming and Optimal Control*. Belmont, MA: Athena Scientific, 2000.

[BBC11] Dimitris Bertsimas, David B. Brown, and Constantine Caramanis. "Theory and Applications of Robust Optimization". In: *SIAM Review* 53.3 (2011), pp. 464–501. DOI: 10.1137/080734510.

[BL92] Fischer Black and Robert Litterman. "Global Portfolio Optimization". In: *Financial Analysts Journal* 48.5 (1992), pp. 28–43. DOI: 10.2469/faj.v48.n5.28.

[Bor19] Nicola Borri. "Conditional tail-risk in cryptocurrency markets". In: *Journal of Empirical Finance* 50 (2019), pp. 1–19. ISSN: 0927-5398. DOI: https://doi.org/10.1016/j.jempfin.2018.11.002.

[BM19] Alexander Brauneis and Roland Mestel. "Cryptocurrency-portfolios in a mean-variance framework". In: *Finance Research Letters* 28 (2019), pp. 259–264. ISSN: 1544-6123. DOI: https://doi.org/10.1016/j.frl.2018.05.008.

[Bra+21] Alexander Brauneis, Roland Mestel, Ryan Riordan, and Erik Theissen. "How to measure the liquidity of cryptocurrency markets?" In: *Journal of Banking and Finance* 124 (2021), p. 106041. ISSN: 0378-4266. DOI: 10.1016/j.jbankfin.2020.106041.

[BL24] Alessio Brini and Jimmie Lenz. "A Comparison of Cryptocurrency Volatility-benchmarking New and Mature Asset Classes". In: *Financial Innovation* 10.122 (2024). DOI: 10.1186/s40854-024-00646-y.

[BF85] Gary P. Brinson and Nimrod Fachler. "Measuring Non-U.S. Equity Portfolio Performance". In: *The Journal of Portfolio Management* 11.3 (1985). Spring, pp. 73–76. DOI: 10.3905/jpm.1985.409005. URL: https://www.pm-research.com/content/iijpormgmt/11/3/73.

[Çel+21] Eranda Çela, Stephan Hafner, Roland Mestel, and Ulrich Pferschy. "Mean-variance portfolio optimization based on ordinal information". In: *Journal of Banking & Finance* 122 (2021), p. 105989. DOI: 10.1016/j.jbankfin.2020.105989.

[Çel+25] Eranda Çela, Stephan Hafner, Roland Mestel, and Ulrich Pferschy. "Integrating multiple sources of ordinal information in portfolio optimization". In: *Annals of Operations Research* 346.3 (2025), pp. 1967–1995. DOI: 10.1007/s10479-025-06495-x.

[Con+23] Lin William Cong, Xi Li, Ke Tang, and Yang Yang. "Crypto Wash Trading". In: *Management Science* 69.11 (Sept. 19, 2023), pp. 6427–6454. DOI: 10.1287/mnsc.2021.02709.

[Cor+20] Shaen Corbet, Charles Larkin, Brian M. Lucey, Andrew Meegan, and Larisa Yarovaya. "The Impact of Macroeconomic News on Bitcoin Returns". In: *The European Journal of Finance* 26.14 (2020), pp. 1396–1416. DOI: 10.1080/1351847X.2020.1737168. URL: https://doi.org/10.1080/1351847X.2020.1737168.

[ČTŽ22] Maria Čuljak, Bojan Tomić, and Saša Žiković. "Benefits of sectoral cryptocurrency portfolio optimization". In: *Research in International Business and Finance* 60 (Feb. 2022). Available online 8 January 2022, p. 101615. ISSN: 0275-5319. DOI: 10.1016/j.ribaf.2022.101615.

[Fab+07] Frank J. Fabozzi, Petter N. Kolm, Dessislava A. Pachamanova, and Sergio M. Focardi. *Robust Portfolio Optimization and Management*. Hoboken, NJ: John Wiley & Sons, 2007. ISBN: 978-0471921226.

[FM24] Angelo Forino and Giacomo Morelli. *Liquidity and Volatility Nexus in Cryptocurrency Markets*. SSRN Working Paper. Sept. 2024.

[FTS] FTSE Russell. *FTSE Grayscale Crypto Sector Indices*. London Stock Exchange Group. URL: https://www.lseg.com/en/ftse-russell/indices/grayscale-crypto-sector (visited on 10/11/2025).

[Gha21] Mahdi Ghaemi Asl. "Good and Bad Volatility Spillovers in the Cryptocurrency Market: New Evidence from a TVP-VAR Asymmetric Connectedness Approach". In: *SSRN Electronic Journal* (2021). URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3957317.

[GK18] Konstantinos Gkillas and Paraskevi Katsiampa. "An application of extreme value theory to cryptocurrencies' returns: Investment risk and value-at-risk forecasting". In: *Financial Innovation* 4.1 (2018), pp. 1–18. URL: https://link.springer.com/article/10.1186/s40854-018-0096-3.

[Gol] Boris Golden. *4Ms: A Framework to Assess Early-Stage Startups*. Slideshare infographic: Market, Model, Management, Momentum. Partech. URL: https://www.slideshare.net/slideshow/4ms-a-framework-for-earlystage-startups/72272649 (visited on 11/11/2025).

[Gue+19] Khaled Guesmi, Samir Saadi, Ilyes Abid, and Zied Ftiti. "Portfolio diversification with virtual currency: Evidence from bitcoin". In: *International Review of Financial Analysis* 63.C (2019), pp. 431–437. DOI: 10.1016/j.irfa.2018.03.004.

[HV21] Daniela Hanicova and Radovan Vojtko. "Rebalancing Premium in Cryptocurrencies". In: *SSRN Electronic Journal* (2021). SSRN Working Paper No. 3982120. URL: https://ssrn.com/abstract=3982120.

[Har69] Frank Harary. *Graph Theory*. Reading, MA: Addison-Wesley, 1969. ISBN: 9780201027877.

[HL99] Guangliang He and Robert Litterman. *The Intuition Behind Black–Litterman Model Portfolios*. Tech. rep. Working Paper. Goldman Sachs Asset Management, 1999.

[HH12] Helge Holden and Lars Holden. "Optimal Rebalancing of Portfolios with Transaction Costs". In: *Stochastics* 85.3 (2012), pp. 371–394. DOI: 10.1080/17442508.2011.651219.

[JL17] Zhengyao Jiang and Jinjun Liang. "Cryptocurrency Portfolio Management with Deep Reinforcement Learning". In: *Proceedings of the 2017 Intelligent Systems Conference (IntelliSys)*. IEEE, Oct. 2017, pp. 905–913. DOI: 10.1109/IntelliSys.2017.8324237. URL: https://ieeexplore.ieee.org/document/8324237.

[Kat17] Paraskevi Katsiampa. "Volatility of Bitcoin: a Markov-switching GARCH approach". In: *Finance Research Letters* 27 (2017), pp. 32–36. URL: https://linkinghub.elsevier.com/retrieve/pii/S1544612317303008.

[KD20] Olivier Kraaijeveld and Johannes De Smedt. "The predictive power of public Twitter sentiment for cryptocurrency returns". In: *Journal of International Financial Markets, Institutions and Money* 65 (2020), p. 101188. DOI: 10.1016/j.intfin.2020.101188.

[La +21] Massimo La Morgia, Alessandro Mei, Francesco Sassi, and Julinda Stefa. "The Doge of Wall Street: Analysis and Detection of Pump and Dump Cryptocurrency Manipulations". In: *arXiv preprint arXiv:2105.00733* (2021). URL: https://arxiv.org/abs/2105.00733.

[Lel99] Hayne E. Leland. *Optimal Portfolio Management with Transactions Costs and Capital Gains Taxes*. Working Paper RPF-290. Institute of Business and Economic Research, University of California, Berkeley, 1999. DOI: 10.2139/ssrn.206871.

[Let+23] Blanka Let, Konrad Sobański, Wojciech Świder, and Katarzyna Włosik. "Investor Sentiment and Efficiency of the Cryptocurrency Market: The Case of the Crypto Fear & Greed Index". In: *Eurasian Business and Economics Perspectives*. Vol. 25. Eurasian Studies in Business and

Economics. Cham: Springer, Oct. 14, 2023, pp. 271–287. DOI: 10.1007/978-3-031-36286-6_16. URL: https://link.springer.com/chapter/10.1007/978-3-031-36286-6_16.

[LJG24] Jian Liu, Jiansuer Julaiti, and Shangde Gou. "Decomposing Interconnectedness: A Study of Cryptocurrency Spillover Effects in Global Financial Markets". In: *Finance Research Letters* 61 (2024), p. 104950. DOI: 10.1016/j.frl.2023.104950. URL: https://doi.org/10.1016/j.frl.2023.104950.

[Liu19] Weiyi Liu. "Portfolio diversification across cryptocurrencies". In: *Finance Research Letters* 29 (2019), pp. 200–205. ISSN: 1544-6123. DOI: https://doi.org/10.1016/j.frl.2018.07.010.

[Ma+20] Yechi Ma, Ferhana Ahmad, Miao Liu, and Zilong Wang. "Portfolio optimization in the era of digital financialization using cryptocurrencies". In: *Technological Forecasting and Social Change* 161 (2020), p. 120265. ISSN: 0040-1625. DOI: https://doi.org/10.1016/j.techfore.2020.120265.

[Mar52] Harry Markowitz. "Portfolio Selection". In: *The Journal of Finance* 7.1 (1952), pp. 77–91. DOI: 10.2307/2975974.

[Meu08] Attilio Meucci. "The Black–Litterman Approach: Original Model and Extensions". In: *SSRN Electronic Journal* (2008). Available at SSRN. DOI: 10.2139/ssrn.1117574.

[MG22] Ivan Mužić and Ivan Gržeta. "Expectations of Macroeconomic News Announcements: Bitcoin vs. Traditional Assets". In: *Risks* 10.6 (June 12, 2022), p. 123. DOI: 10.3390/risks10060123. URL: https://www.mdpi.com/2227-9091/10/6/123.

[PS88] André F. Perold and William F. Sharpe. "Dynamic Strategies for Asset Allocation". In: *Financial Analysts Journal* 44.1 (1988), pp. 16–27. DOI: 10.2469/faj.v44.n1.16.

[21] *Prudential treatment of cryptoasset exposures. Consultative Document.* Consultative document d519. Issued for comment by 10 September 2021. Basel: Basel Committee on Banking Supervision, June 2021, p. 25. URL: https://www.bis.org/bcbs/publ/d519.pdf.

[SP] S&P Dow Jones Indices. *S&P Cryptocurrency DeFi Top 15 Index.* S&P Dow Jones Indices — S&P Global. URL: https://www.spglobal.com/spdji/en/indices/digital-assets/sp-cryptocurrency-defi-top-15-index (visited on 10/11/2025).

[Sal24] A. Salo. "Fifty Years of Portfolio Optimization". In: *European Journal of Operational Research* 319.2 (2024), pp. 369–384. DOI: 10.1016/j.ejor.2023.00327.

[Sch24] Stefan Scharnowski. "Dark Web Traffic, Privacy Coins, and Cryptocurrency Trading Activity". In: *Finance Research Letters* 67 (2024), p. 105875. DOI: 10.1016/j.frl.2024.105875. URL: https://doi.org/10.1016/j.frl.2024.105875.

[Sim25] Simmons & Simmons LLP. *Tax on Cryptocurrency.* Accessed: 10 Oct. 2025. 2025. URL: https://www.simmons-simmons.com/en/features/tax-on-cryptocurrency (visited on 10/10/2025).

[Sin21] Kevin Singpurwala. "Sentiment Analysis Trading Indicators: Exemplified by Natural Language Processing and Bitcoin". In: *Proceedings of the 35th Twente Student Conference on IT (TSConIT 2021).* University of Twente, Faculty of Electrical Engineering, Mathematics and Computer Science. Enschede, The Netherlands, 2021. URL: https://essay.utwente.nl/fileshare/file/87003/tosubmitwith_corrections%20%281%29.pdf.

[SSN24] Sutta Sornmayura, Nichanan Sakolvieng, and Kaimook Numgaroonaroonroj. "Optimizing Cryptocurrency Portfolios: A Comparative Study of Rebalancing Strategies". In: *Journal of Finance and Business Research* 8.4(1) (2024). DOI: 10.35609/jfbr.2024.8.4(1).

[Thé25] Josué Thélissaint. *Assessing Cryptomarket Risks: Macroeconomic Forces, Market Shocks and Behavioural Dynamics.* SSRN working paper; last revised 2025-04-08. Mar. 31, 2025. DOI: 10.2139/ssrn.5075457. URL: https://ssrn.com/abstract=5075457.

[Tuk77] John W. Tukey. *Exploratory Data Analysis.* Internet Archive edition. Reading, MA: Addison-Wesley Publishing Company, 1977. ISBN: 9780201076165. URL: https://archive.org/details/exploratorydataa00tuke_0.

[Zha+22] Yu Zhang, Harshdeep Ahluwalia, Allison Ying, Michael Rabinovich, and Aidan Geysen. *Rational Rebalancing: An Analytical Approach to Multiasset Portfolio Rebalancing Decisions and Insights.* Vanguard Research Report. Accessed: Oct. 10, 2025. The Vanguard Group, Oct. 2022. URL: https://corporate.vanguard.com/content/dam/corp/research/pdf/rational_rebalancing_analytical_approach_to_multiasset_portfolio_rebalancing.pdf.

A Fear and Greed Construction Comparison

Table 4: **Side-by-side overview of Fear & Greed index constructions**

	F&G Crypto markets		F&G Traditional markets	
Item	**Weight**	**Description**	**Weight**	**Description**
Volatility	25%	Current volatility and maximum drawdowns of Bitcoin compared to the corresponding average values over the last 30 and 90 days.	14.3%	VIX index versus 50-day moving average.
Market momentum and volume	25%	Approximated by current market volume (buy/sell) and momentum (compared to averages over the last 30/90 days); the two values are stacked together.	14.3%	Stock Price Momentum: S&P 500 vs. 125-day moving average.
			14.3%	Stock Price Strength: number of 52-week highs vs. lows.
			14.3%	Stock Price Breadth: volume in rising vs. declining stocks.
Social media	10%	For each coin, posts containing various hashtags on Twitter/X are aggregated and counted; interaction speed and counts over specific windows proxy public interest (higher interaction $\Rightarrow$ greedier market).	0%	None.
Dominance metrics	10%	Bitcoin's share of total crypto market capitalization; rising BTC dominance with declining alt-coin activity is interpreted as a sign of fear.	0%	None.
Trends	10%	Google Trends data for Bitcoin-related search queries.	14.3%	Put/Call options: option-volume balance (bullish/bearish sentiment).
Surveys	15%	Based on a large public polling platform, including several questions about respondents' views on the crypto market.	14.3%	Safe-haven demand: stocks vs. Treasury bonds.
			14.3%	High-yield demand: HY–IG spread.

Note: Frequency observations daily. Final score averages these indicators and ranges 0 (extreme fear) to 100 (extreme greed). The seven traditional Fear and Greed components are equally weighted.

B Derivations for the Portfolio Construction Engine

Investor views as linear constraints. Let $r \in \mathbb{R}^N$ be the vector of (excess) asset returns. A set of K views is encoded by a pick matrix $P \in \mathbb{R}^{K\times N}$ and a view vector $q \in \mathbb{R}^K$ such that $P\,r = q$. The diagonal matrix $\Omega \in \mathbb{R}^{K\times K}$ contains the variance (confidence) of each view. This formulation accommodates absolute, relative, and ordinal views.

Equilibrium prior returns. Given market weights $w^{\text{mkt}} \in \mathbb{R}^N$, covariance matrix Σ, and risk-aversion δ, the implied equilibrium excess returns are

$$\pi \;=\; \delta\,\Sigma\,w^{\text{mkt}},$$

where δ is often estimated as $\delta = (\mathbb{E}[R^m] - r_f)/\sigma_m^2$.

Bayesian update (posterior returns). Define $\tau > 0$ to scale the prior uncertainty. The posterior mean μ solves the mean–variance view optimization

$$\mu \;=\; \arg\min_r \;\; (r-\pi)^\top(\tau\Sigma)^{-1}(r-\pi) \;+\; (q-Pr)^\top\Omega^{-1}(q-Pr),$$

with closed-form solution

$$\mu \;=\; \left[(\tau\Sigma)^{-1} + P^\top\Omega^{-1}P\right]^{-1}\Big[(\tau\Sigma)^{-1}\pi + P^\top\Omega^{-1}q\Big].$$

This is the standard Black–Litterman formula.

Extensions via optimization. When views are ordinal or inequality-based, one modifies the optimization problem by adding linear constraints (e.g., $P_1 r \geq P_2 r$) or penalty terms. Additional regularization can enforce bounds on μ or shrink deviations from the prior. These extensions preserve convexity and allow robust or distributionally-aware variants.

Robust mean–variance optimization. Once μ is obtained, we solve for weights $w \in \mathbb{R}^N$:

$$\min_{w} \ \tfrac{1}{2} w^\top \Sigma\, w \ - \ \lambda\, w^\top \mu \ + \ \lambda\, u \sqrt{w^\top E\, w},$$

subject to long-only, full-investment and concentration constraints. Here $E \succeq 0$ is the covariance of the estimation error in μ, and $u \geq 0$ controls . This second-order cone program yields weights that are less sensitive to forecasting error while achieving the desired risk–return trade-off.

Implementation notes. In our implementation the statistics (Σ, π) and signals are updated weekly using a 60-day lookback window; signals are time-decayed so that recent information receives higher weight. Practitioners may substitute their own baseline returns or weights, or incorporate additional constraints (e.g., sector caps, minimum weight floors). All such modifications fit within the BL framework described above.