Modern Portfolio Theory in the Crypto-Wilderness

Ivan Vynyavskyy
Complexity Science Hub & TU Wien, Austria

Stefan Kitzler
Complexity Science Hub & Austrian Institute of Technology, Austria

Bernhard Haslhofer
Complexity Science Hub, Austria

Aviv Yaish
Yale University & IC3, USA
Complexity Science Hub, Austria

Abstract

Modern portfolio theory (MPT) prescribes how to maximise the return of an asset portfolio for a given level of risk. The optimal trade-off between return and variance defines the efficient frontier. Whether actual cryptoasset portfolios approximate this prescription and whether proximity to the frontier translates into realised performance remain difficult to test at large scale in traditional markets due to their opaque nature and the inaccessibility of data. As we show, public blockchains make these questions measurable: every token transfer is recorded, thus enabling complete portfolio reconstruction for every account at any point in time. We leverage this transparency to reconstruct cryptoasset portfolios for over 116 M Ethereum accounts across the full chain history (2015–2025), measure their distance to the constrained efficient frontier, and quantify how deviations translate into realised performance. Here we show that market entry timing, not allocation choice, is the dominant predictor of realised cryptoasset returns. On-chain wealth is highly concentrated and portfolios are pervasively under-diversified, with single-asset holdings accounting for 83.35% of accounts. Two-asset portfolios sit closest to the efficient frontier defined by their held assets, a proximity that reflects the narrowness of their opportunity set rather than deliberate optimisation. Passive market-capitalisation weighting outperforms every MPT optimisation strategy in median realised return, and entry month alone explains 70–79% of the variance in returns, far exceeding the contribution of allocation choice. Mean-variance optimisation therefore appears neither descriptive of observed behaviour nor prescriptively useful in the cryptoasset domain, even if MPT retains its value as a normative benchmark.

2012 ACM Subject Classification Applied computing → Digital cash; Mathematics of computing → Mathematical optimization; Applied computing → Economics

Keywords and phrases Modern Portfolio Theory, Mean-Variance Optimisation, DeFi, Blockchain

Acknowledgements This work is partially funded by the Austrian security research program KIRAS of the Federal Ministry of Finance (BMF) under the project LLEA (grant agreement 926183). The computational results have been achieved using the Austrian Scientific Computing (ASC) infrastructure.

1 Introduction

Modern portfolio theory (MPT), for which Markowitz received the Nobel Prize in Economics [34], prescribes that portfolios should lie on the mean-variance efficient frontier, that is, maximise expected return for a given level of risk. MPT's core assumptions are that investors are *perfectly rational*, actively optimising their allocation, and *risk-averse*, preferring lower variance for equal return. These assumptions are idealisations challenged in this paper.

Cryptoasset markets are a uniquely suitable empirical test bed for MPT. First, blockchains record every token transfer in a publicly auditable ledger, enabling the reconstruction of complete portfolio histories for every account, unlike traditional markets where holdings are observable only through selective brokerage disclosures. Second, tokens are accessible permissionlessly and adhere to standards such as ERC-20 so on-chain portfolio choices reflect actual allocation decisions rather than institutional constraints. Third, the Ethereum ecosystem is particularly well-suited for this examination, spanning a decade of multiple market cycles and hosting the largest decentralised finance ecosystem. This setting enables us to examine both assumptions empirically at population scale.

In traditional markets, investors are systematically under-diversified and often underperform passive benchmarks [20, 8, 1], while on-chain studies reveal highly concentrated token distributions and centralised governance [46, 37]. Existing crypto-portfolio work applies MPT only to hypothetical asset baskets [6, 41, 44, 40] rather than to actual on-chain holdings, and each strand relies on sampled or selected data over limited windows. Yet no study has measured, at scale, how actual on-chain portfolios compare to MPT-optimal allocations, or whether deviating from the constrained efficient frontier has performance consequences.

We fill this gap with the first large-scale empirical study comparing *actual* on-chain cryptoasset portfolios to rebalanced *theoretical optimal* allocations under MPT. We reconstruct on-chain cryptoasset portfolios at scale and pursue three objectives: First (**Obj A**), we characterise the composition and diversification patterns of actual on-chain cryptoasset portfolios. Second (**Obj B**), we quantify the deviation of actual portfolios from the mean-variance efficient frontier under optimisation strategies and held-asset constraints. Third (**Obj C**), we examine how such deviations affect portfolio performance relative to the market.

To summarise, we contribute the following:

1. **Portfolio reconstruction method.** We propose a method for reconstructing actual on-chain token portfolios from transfer events, processing 4562 tokens across the full Ethereum history (2015–2025) and covering 116 M accounts via an incremental ledger-based approach validated against on-chain balances (Section 3). We contribute to open science by making the code (and the full dataset upon publication) publicly available[1].
2. **Actual portfolio composition and diversification.** Aggregate on-chain wealth peaks at \$1.47 T, but is highly concentrated: the vast majority of accounts hold less than \$100, and 99.03% of token wealth is controlled by the top 1% of holders, predominantly centralised cryptoasset exchanges and decentralised finance (DeFi) protocols. Single-asset portfolios account for 83.35% of all accounts, and only 2.17% hold five or more assets. Multi-asset portfolios become more prevalent as account wealth increases, a pattern inconsistent with rational, risk-averse portfolio optimisation (Section 4).
3. **Deviation from efficient frontier.** We measure the ℓ_1 weight distance between actual portfolios and three frontier strategies (minimum variance, maximum return, maximum Sharpe ratio). Portfolio size and suboptimality follow a power-law decay: small portfolios (< 5 assets) lie close to the efficient frontier by construction, while larger portfolios deviate substantially, with distances reaching 48.5%–70.2% depending on the strategy (Section 5).
4. **Deviation effects.** We compare realised returns of actual portfolios against MPT-optimal counterfactuals. Proximity to the constrained efficient frontier does not predict better risk-adjusted performance. Instead, market entry timing, rather than allocation, is the dominant predictive feature of realised returns (Section 6).

[1] `https://github.com/ivan-vynyavskyy/modern-portfolio-theory-in-the-crypto-wilderness`

The results have implications for both theory and practice. For theory, pervasive under-diversification and the dominance of market entry timing over frontier proximity challenge the descriptive adequacy of MPT in the cryptoasset domain. The mismatch points to behavioural factors and a dominant common market factor, rather than mean-variance optimisation, as the organising principles of actual cryptoasset portfolios.

For practice, our findings caution against importing mean-variance frameworks wholesale into cryptoasset portfolio construction: rebalancing toward MPT optima does not reliably improve risk-adjusted performance, and passive market-capitalisation weighting offers a more robust default. The persistence of single-asset portfolios and the extreme wealth concentration we document further suggest that frictions, behavioural factors, and investor-protection concerns deserve greater attention than allocation theory in shaping interventions across retail, advisory, platform, and regulatory contexts.

The remainder of the paper reviews related work, presents the method and findings for each objective, and concludes with a discussion of implications and limitations.

2 Background and Related Work

This section reviews the MPT framework and efficient frontier, empirical portfolio studies in traditional finance, organised around our three objectives, empirical studies in blockchain markets, and closes with the synthesis and research gap that motivates the present work.

2.1 Modern Portfolio Theory

The MPT framework was introduced by Markowitz [33], who proposed a method for allocating assets to maximise expected returns for a given level of risk, known as mean-variance optimisation. MPT relies on several assumptions: investors are perfectly rational and risk-averse; markets are frictionless (no transaction costs, taxes, or borrowing constraints); and prices are efficient, instantaneously reflecting all available information. Building on MPT, Sharpe [43] introduced the capital asset pricing model (CAPM), which models an asset's expected return as a function of its systematic risk relative to the market.

Central to MPT is the concept of the *efficient frontier*, the upper boundary of the minimum-variance set in expected-return/standard-deviation space: portfolios on the frontier achieve the highest possible expected return for a given level of risk, while portfolios below it are suboptimal. Any actual portfolio can thus be evaluated against the frontier, quantifying the gap between what an account holds and what mean-variance optimisation prescribes. For a comprehensive introduction to MPT and relevant literature, see Elton et al. [16].

2.2 Empirical Studies in Traditional Finance

A large body of empirical research has studied the portfolios of investors in traditional assets. Campbell [9] establishes household finance as a field, surveying evidence that while many households invest effectively, a significant minority (typically poorer and less educated) make portfolio mistakes. The findings summarised below provide context for each of our objectives.

Composition and diversification. A central finding is that investors hold fewer assets than theory prescribes. Goetzmann and Kumar [20] analyse more than 40,000 equity accounts at a U.S. discount brokerage and find that the average investor holds only four stocks (median of three), with greater under-diversification among younger, low-income, and less-educated investors. This concentration often takes the form of familiarity bias: French and Poterba [19] document strikingly high domestic ownership shares despite substantial

diversification gains from international investing, and Huberman [24] shows that shareholders concentrate holdings in geographically familiar companies. Benartzi and Thaler [3] further show that 401(k) participants tend to divide contributions equally across available options, the so-called *1/n heuristic*, another departure from mean-variance optimisation; DeMiguel et al. [14] find that this naive rule is hard to beat by mean-variance strategies out-of-sample, while Hellum et al. [23] show that this gap closes once the number of assets exceeds the number of training observations. Together, these studies characterise actual portfolios as simpler and less diversified than MPT prescribes.

Deviation from optimal allocations. Quantifying how far actual portfolios lie from theoretical optima requires both a benchmark and a distance measure. Using comprehensive administrative data on all Swedish households, Calvet et al. [8] measure the welfare costs of investment mistakes: while most households actually outperform their domestic stock index through international diversification, a minority incurs substantial losses from under-diversification. Massa and Simonov [35] further show that Swedish investors concentrate holdings in stocks related to their employment income and geographic proximity rather than hedging against non-financial income risk, a systematic deviation from theoretically optimal allocations. At the fund level, Didier et al. [15] find that even mutual funds exhibit limited international diversification, with global funds potentially achieving 2.6–5.5% higher annual risk-adjusted returns by broadening their allocations. These studies typically rely on welfare proxies or return differentials rather than portfolio-level distance metrics, leaving open how directly observable allocations compare to the efficient frontier.

Performance consequences. Individual investors generally underperform passive benchmarks. Barber and Odean [1] find that among 66,465 household accounts, the most active traders earned 11.4% annually versus 17.9% for the market, with overconfidence driving excessive trading. Odean [38] documents the disposition effect: investors realise gains more readily than losses, reducing annual returns by 4.4 percentage points. Classically, Brinson et al. [7] show that asset allocation policy explains 93.6% of the variation in quarterly portfolio returns across 91 large U.S. pension funds, with security selection and market timing playing minor roles, a benchmark our cryptoasset findings invert.

2.3 Empirical Studies in Blockchain Markets

Blockchain ledgers provide a uniquely complete empirical record of asset holdings; this subsection describes the Ethereum setting that makes portfolio reconstruction possible and reviews prior measurement work on token distributions and crypto-portfolio optimisation.

Ethereum as an empirical setting. Fungible tokens (ERC-20) on Ethereum are digital assets managed by smart contracts (contract accounts (CAs)) that track balances for each account. Ethereum distinguishes two account types: externally owned accounts (EOAs), controlled by private keys, hold tokens directly; CAs implement programmable logic and hold tokens on behalf of protocols such as decentralised finance (DeFi) applications for lending, trading, and liquidity provision. Every token transfer is recorded as an event in the blockchain logs, providing a complete and auditable history of all token movements and forming the empirical basis for portfolio reconstruction from public data.

On-chain measurement studies. Prior work has leveraged this data to characterise token distributions and network structure. Nadler and Schär [37] analyse the token distribution of DeFi governance tokens and find that protocol governance is heavily centralised. Victor and Lüders [46] conduct an early empirical study of Ethereum ERC-20 tokens over the first 6 million blocks, analysing token distributions and trade networks. These studies establish that on-chain holdings are concentrated, but stop short of reconstructing portfolios

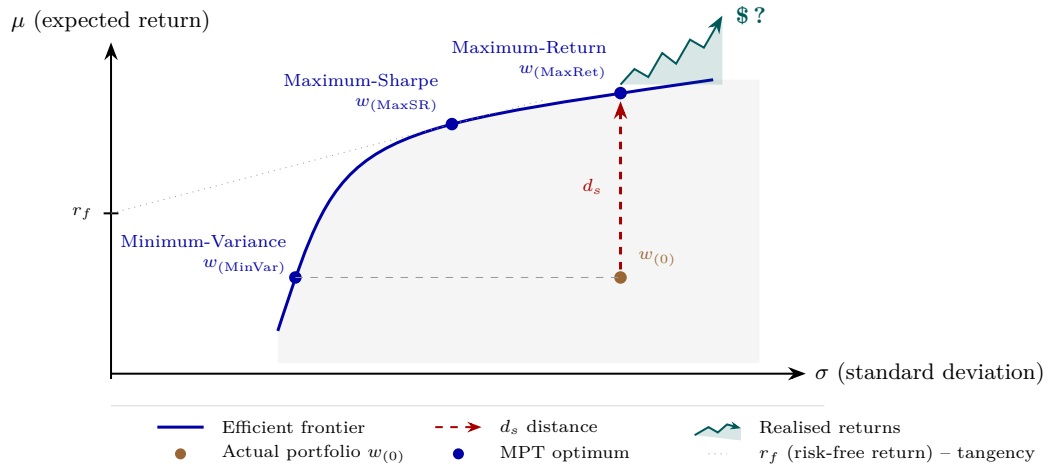

Figure 1 Overall analytical approach in expected-return/standard-deviation space. Each actual portfolio $\bullet\, w_{(0)}$ is characterised (**Obj A**), compared to three frontier optima $\bullet$ on the efficient frontier — via the weight-space distance $\dashrightarrow d_s$ (**Obj B**), and evaluated on realised returns (**Obj C**).

at the account level or evaluating them against a portfolio-theoretic benchmark.

MPT applied to cryptoassets. A separate strand applies MPT to cryptoassets, typically on curated baskets rather than actual holdings. Brière et al. [6] examine Bitcoin's role in diversified portfolios; Platanakis and Urquhart [41], Trimborn et al. [44], and Petukhina et al. [40] apply mean-variance optimisation to small baskets of major cryptoassets. Borri et al. [5] apply an empirical asset pricing approach to tokens and find that their risk-adjusted performance is comparable to traditional markets. None of these studies compares actual account holdings to MPT-optimal allocations at population scale.

2.4 Synthesis and Research Gap

Together, these three strands establish that actual investors in traditional markets systematically deviate from MPT prescriptions, that on-chain portfolios can be measured at scale from public blockchain data, and that prior work on crypto-portfolio optimisation has focused on hypothetical baskets rather than actual holdings. Yet no study has characterised actual on-chain portfolio composition at population scale, measured how these portfolios compare to MPT-optimal allocations, or tested whether deviating from the efficient frontier has measurable performance consequences. Blockchain data enables this analysis at a completeness and scale impossible with traditional brokerage data; the present work fills this gap.

3 Data and Methods

3.1 Overall Approach

Our analysis reconstructs actual on-chain portfolios and evaluates them against MPT-optimal allocations along the three objectives. Figure 1 provides a schematic overview.

Each actual portfolio is represented by its asset-weight vector $w_{(0)}$, corresponding to expected return $\mu_{(0)}$ and standard deviation $\sigma_{(0)}$. We place $w_{(0)}$ in the feasible (μ, σ) region and characterise its composition and diversification (Obj A). We compute its weight-space distance d_s to three reference portfolios on the efficient frontier: the minimum-variance (MinVar), maximum-return (MaxRet), and maximum-Sharpe-ratio (MaxSR) optima. This

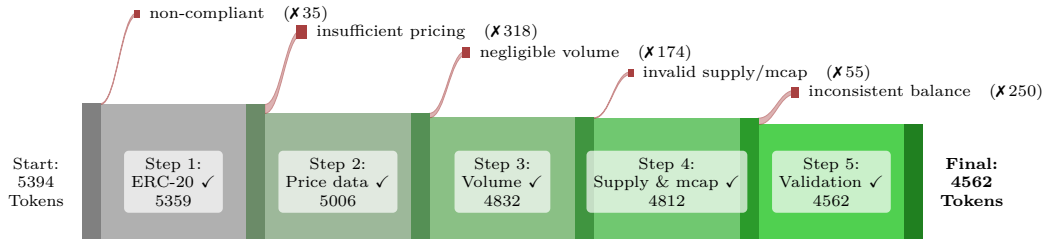

Figure 2 Token cleaning pipeline. Out of 5394 initial tokens, 4562 pass all filtering stages (ERC-20 compliance; price history, trading volume, and supply/market-cap validation; reconstruction validation) and are retained for subsequent analysis.

quantifies deviation from MPT prescriptions (Obj B). We compare the realised returns of $w_{(0)}$ against those references over matched holding periods, testing whether proximity to the frontier predicts risk-adjusted performance (Obj C).

The remainder of this section describes the data sources underlying our analysis, the method for reconstructing on-chain portfolios from transfer events, the computation of MPT-optimal reference portfolios, and the distance and performance metrics used to evaluate actual portfolios against these references.

3.2 Data Collection

We collect data from two sources: the Ethereum blockchain, which provides transactions and event logs, and the Coingecko platform [11], which provides token metadata and prices.

- **Ethereum blockchain:** We run a full node [17] and extract all transactions and event logs from the genesis block (30 July 2015, block #0) to 5 December 2025 (block #23,943,069). Smart contracts emit information to the blockchain *logs* when on-chain events occur. In particular, `Transfer` events following the ERC-20 standard record token movements between accounts and form the empirical basis for portfolio reconstruction.
- **Coingecko:** We collect metadata for the top tokens on Coingecko, ranked by market capitalisation, including the Ethereum contract address, name, and symbol. For each token we additionally collect daily USD closing prices, market capitalisation, fully diluted valuation (FDV = Price × Total Supply), total circulating supply, and trading volume.

Figure 2 summarises the multi-stage cleaning pipeline described below. Applied to the 5394 initial tokens, 4562 (≈84.5%) pass all stages and are retained for subsequent analysis.

1. **Standard compatibility.** We verify ERC-20 compliance by extracting the runtime bytecode from `CREATE` and `CREATE2` traces and checking for ERC-20 function and event selectors [46]. Proxy contracts are resolved to their underlying implementation.
2. **Market activity.** We apply three filters for sufficient trading activity and reliable price data. First, we exclude tokens with fewer than 15 days of price history, which is insufficient for return estimation. Second, we exclude tokens with negligible historical trading volume. Third, we validate supply data, excluding tokens whose market capitalisation or FDV exceeds Ethereum's (an indicator of misconfigured supply).
3. **On-chain consistency.** Finally, we validate that our event-based reconstruction matches the on-chain state. For each candidate token, we reconstruct balances at randomly sampled (account, block) pairs and compare them to the on-chain balance retrieved via RPC. Tokens with any disagreement between reconstructed and on-chain balances are excluded,

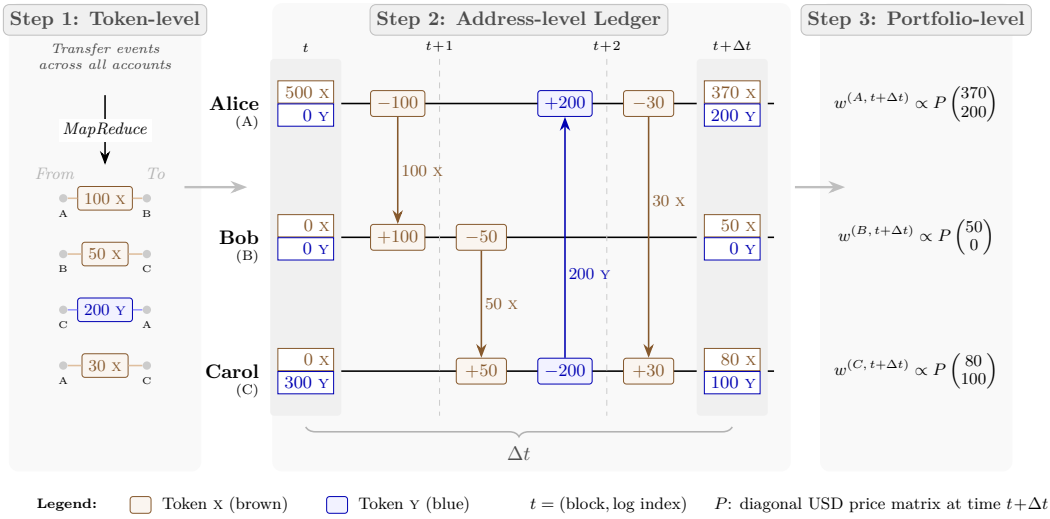

Figure 3 Three-step portfolio reconstruction pipeline for two example tokens (X, brown; Y, blue) and three example accounts (Alice, Bob, Carol). **Step 1** (left): a MapReduce procedure extracts `Transfer` events from blockchain logs and produces per-token transfer files. **Step 2** (centre): for each token, transfers are recorded as signed ledger entries (debit − for senders, credit + for receivers), ordered by block height and log index. Cumulative sums over the interval Δt yield per-account balances at the snapshot time t. **Step 3** (right): per-account balances at snapshot t are combined with the asset prices $p_i^{(t)}$ to yield position values $v_i^{(a,t)}$, which are normalised into portfolio weight vectors $w^{(a,t)}$. In our analysis, Δt corresponds to one month.

since discrepancies indicate non-standard token implementations (e.g., rebasing tokens, tokens with transfer hooks, or tokens that modify balances outside `Transfer` events).

3.3 Portfolio Reconstruction

As Figure 3 shows, starting from raw Ethereum *logs* and *traces*, we reconstruct the token holdings of each account at any point in time (block height) in three steps: extracting token-level transfer events, aggregating transfers into per-address balance ledgers, and computing portfolio snapshots at target time points. Implementation details on the underlying storage infrastructure are provided in Appendix E.

Step 1: Token-level extraction. We extract the logs for the ERC-20 transfer-signature topic and write per-token files of (sender, recipient, value) tuples, parallelised across tokens via a MapReduce pipeline [12]. Some assets require additional token-specific handling: Wrapped ETH (WETH), for instance, exposes `Deposit` and `Withdrawal` events, which we map to synthetic `Transfer` entries from or to the zero address to unify the event stream.

Step 2: Account-level ledger. We transform the `Transfer` stream into per-(account, token) ledgers, where each account corresponds to an Ethereum address. Each event produces a signed entry (debit for sender, credit for recipient), ordered by block height and log index. The ledger supports balance aggregation over any block range from genesis to a target height.

Step 3: Portfolio snapshots. A portfolio is the collection of tokens held by an account at a given point in time, with Ethereum block heights converted to datetimes via an internal mapping. For each account and target timestamp, we sum the corresponding ledger entries, price each position with forward-filled daily USD closes, and normalise the position values into a fully-invested weight vector. We sample *monthly snapshots* at the first day of each

calendar month, matching the literature [20] and balancing resolution against compute.

For each account a and time t, we obtain the USD value of each held position $v_i^{(a,t)}$ for asset i, the total portfolio value $V^{(a,t)} = \sum_i v_i^{(a,t)}$, the number of held tokens $N^{(a,t)}$, and the fully-invested weight vector[2]

$$w^{(a,t)} = \left(w_1^{(a,t)}, \ldots, w_N^{(a,t)}\right), \qquad w_i^{(a,t)} = \frac{v_i^{(a,t)}}{V^{(a,t)}}, \tag{1}$$

where $\sum_{i=1}^{N} w_i^{(a,t)} = 1$. Reconstruction correctness is validated against on-chain balances as described in Section 3.2 (step 3).

3.4 Efficient Frontier Computation

Building on the MPT framework introduced in Section 2.1, we estimate the efficient frontier for each account and snapshot date t from empirical return statistics over a historical lookback window, constrained to the account's already-held assets. Token-specific functionalities (e.g., underlying lending positions, staking yields) are disregarded: tokens are treated as simple assets characterised by prices.

Return and risk estimation.

▶ **Definition 1** (Expected return). *The expected asset-return of asset i over a $T^{(-)}$-day lookback window ending at t is*

$$\mu_i^{(t)} \;=\; \frac{1}{T^{(-)}} \sum_{\tau=0}^{T^{(-)}-1} r_i^{(t-\tau)}, \qquad \textit{with} \qquad r_i^{(t)} \;=\; \log\left(\frac{p_i^{(t)}}{p_i^{(t-1)}}\right)$$

the log-return of asset i at time t and $p_i^{(t)}$ the USD closing price of token i at time t.

▶ **Definition 2** (Portfolio mean and variance). *The expected return and variance of account a's portfolio at time t are*

$$\mu^{(a,t)} \;=\; \sum_{i=1}^{N} w_i^{(a,t)} \mu_i^{(t)}, \qquad \left(\sigma^{(a,t)}\right)^2 \;=\; \left(w^{(a,t)}\right)^{\top} \Sigma^{(t)}\, w^{(a,t)},$$

where $\Sigma^{(t)} \in \mathbb{R}^{N\times N}$ is the empirical covariance matrix over the lookback ending at t, and N is the number of assets held at t.

Reference portfolios under optimisation strategies.

For each account a and time t, we apply three optimisation strategies, $s \in \{\text{MinVar}, \text{MaxRet}, \text{MaxSR}\}$, each producing a reference portfolio $w_{(s)}$ on the constrained efficient frontier.

▶ **Definition 3** (Frontier portfolios). *Let $w_{(0)} \equiv w^{(a,t)}$ denote the actual portfolio with moments $\mu_{(0)}, \sigma_{(0)}$, and $w_{(s)}$ the portfolio optimised under strategy s with moments $\mu_{(s)}, \sigma_{(s)}$:*
Min variance at target return (MinVar): $\min_w \sigma_{(s)}$ *s.t.* $\mu_{(s)} = \mu_{(0)} \;\longrightarrow\; w_{(MinVar)}$.
Max return at target volatility (MaxRet): $\max_w \mu_{(s)}$ *s.t.* $\sigma_{(s)} = \sigma_{(0)} \;\longrightarrow\; w_{(MaxRet)}$.
Max Sharpe ratio (MaxSR): $\max_w (\mu_{(s)} - r_f)/\sigma_{(s)} \;\longrightarrow\; w_{(MaxSR)}$.

---

[2] *Notation convention.* Throughout, when the (a,t) indices are suppressed (e.g. w, V, N, μ, σ), the quantity refers to the portfolio of account a at time t.

The MinVar and MaxRet anchor to the actual portfolio's expected return and standard deviation, respectively, corresponding to the horizontal and vertical projections of $w_{(0)}$ onto the frontier (Figure 1). The MaxSR is the *tangency portfolio*, the point where the line from r_f touches the frontier, and it depends on the assets held but not on their initial weights. We set $r_f = 5\%$ per annum for the MaxSR optimisation, close to on-chain stablecoin lending yields and consistent with [41], who proxy r_f by the U.S. Treasury bill rate (2023–2024 [4]). Section C.2 motivates the calibration and shows that results are insensitive to r_f.

Shared constraints.

All three optimisation problems share the following constraints, which together with the held-asset restriction define the *constrained efficient frontier* used throughout the paper:

- *Fully invested*: $\sum_i w_{(s),i} = 1$, with $w_{(s),i} = 0$ whenever $w_{(0),i} = 0$ (pure reallocation across held assets, no external funds).
- *Long-only*: $w_{(s),i} \geq 0$ (no spot-level short-selling available for most ERC-20 tokens).
- *Single-asset cap*: $w_{(s),i} \leq w_{\max} = 0.9$ (prevents collapse to a near-single-asset solution).

Configuration.

We solve all three optimisation problems with sequential least-squares programming (SLSQP). The following parameters are shared across all optimisation instances:

- *Lookback window:* $T^{(-)} = 60$ daily observations (minimum 45 per asset required, with tokens below the threshold excluded at that time point). The window balances stable covariance estimation against the frequent regime changes of cryptoasset markets.
- *Covariance shrinkage:* We apply Ledoit–Wolf shrinkage [28] to stabilise the empirical covariance matrix when N is large relative to $T^{(-)}$.
- *Return regularisation:* Sample mean returns are highly noisy over the $T^{(-)}$-day lookback in crypto markets. We shrink the sample mean toward a neutral prior of equal expected returns across assets:

  $$\mu_i^{\text{shrunk}} = \lambda\,\bar{\mu} + (1-\lambda)\,\mu_i,$$

  where $\bar{\mu} = \frac{1}{N}\sum_{j=1}^{N}\mu_j$ is the cross-asset mean and $\lambda = 0.5$ (further details in Section C.1).

3.5 Distance and Performance Metrics

We evaluate the optimisation outcomes along two dimensions: how far each optimised portfolio deviates from the actual allocation (Obj B), and whether the resulting shift improves risk-adjusted returns (Obj C).

Weight ℓ_1 distance.

▶ **Definition 4** (Weight ℓ_1 distance)**.** *The ℓ_1 distance between the actual allocation $w_{(0)}$ and an optimised reference $w_{(s)}$ is*

$$d_s = \tfrac{1}{2}\left\|w_{(0)} - w_{(s)}\right\|_1 = \frac{1}{2}\sum_{i=1}^{N}\left|\,w_{(0),i} - w_{(s),i}\,\right|.$$

We use ℓ_1 rather than Euclidean (ℓ_2) because d_s admits a direct financial interpretation: it equals the minimum one-way turnover required to rebalance from $w_{(0)}$ to $w_{(s)}$, and is bounded in $[0, 1]$ for long-only, fully-invested portfolios, with 0 indicating identical allocations and 1 disjoint ones. See Section C.1 for the formal derivation.

Risk-adjusted performance.

▶ **Definition 5** (Forward-realised return). *The forward-realised return of account a's allocation $w^{(a,t)}$ between consecutive snapshots t and $t' = t + T^{(+)}$ is*

$$r_{t'}^{(a,t)} = \sum_{i=1}^{N} w_i^{(a,t)} \left(\frac{p_i^{(t')}}{p_i^{(t)}} - 1 \right),$$

where $p_i^{(t)}$ and $p_i^{(t')}$ are the asset prices at the start and end of the holding window; the expression is evaluated at the actual $w_{(0)}$ or any theoretical $w_{(s)}$.

Raw returns conflate allocation with market movement: in a bull market, even naive allocations are positive. To isolate the allocation's contribution, we apply the capital asset pricing model (CAPM) and decompose the return into a component proportional to the portfolio's market exposure (β) and an excess residual (α), the risk-adjusted return. As the market benchmark $r_{(m)}^{(t)}$, we take the equally-weighted average of the daily log-returns of wrapped ETH (wETH) and wrapped BTC (WBTC). Equal weighting prevents structural overlap with the market-cap-weight strategy.

▶ **Definition 6** (Market exposure). *The market exposure of account a's portfolio at time t is*

$$\beta^{(a,t)} = \sum_i w_i^{(a,t)} \beta_i^{(t)}, \qquad \text{with} \qquad \beta_i^{(t)} = \frac{\mathrm{Cov}_\tau\left(r_i^{(t-\tau)}, r_{(m)}^{(t-\tau)}\right)}{\mathrm{Var}_\tau\left(r_{(m)}^{(t-\tau)}\right)}$$

the sensitivity of asset i to the market at time t, estimated over the same $T^{(-)}$-day lookback, and $w_{(s)}$ admissible in place of $w^{(a,t)}$.

A value $\beta^{(a,t)} = 1$ implies lockstep movement with the market, $\beta^{(a,t)} > 1$ amplifies market moves, and $\beta^{(a,t)} < 1$ dampens them.

▶ **Definition 7** (Risk-adjusted excess return). *The risk-adjusted excess return of account a's portfolio at t' is*

$$\alpha_{t'}^{(a,t)} = r_{t'}^{(a,t)} - \beta^{(a,t)} \cdot r_{(m)}^{(t')}.$$

Here, we set $r_f = 0$ to avoid an off-chain rate assumption. Positive $\alpha_{t'}^{(a,t)}$ indicates outperformance beyond market exposure, and negative values indicate underperformance.

Ethics and data handling.

The study uses public Ethereum logs, public token-price data, and aggregate address tags. We do not deanonymise addresses, and address tags are used only for aggregate contract/category-level analysis. The released artifacts do not expose private address-level information, only public data already publicly available on-chain.

4 Portfolio Characteristics

This section addresses **Obj A** by characterising the wealth, concentration, and diversification of on-chain portfolios. The analysis covers 4562 ERC-20 tokens and 116 182 156 (116 M) accounts at monthly snapshots across Ethereum's history (July 2015 – December 2025).

4.1 Evolution of Wealth

We measure account wealth V as the USD value of assets held at each monthly snapshot. We start with an aggregate view by summing these values across all accounts, providing an overview of the total ERC-20 wealth in the ecosystem and its evolution over time.

Figure 4 traces the total USD value of reconstructed portfolios over time and contrasts it with total value locked (TVL) [13]. Total wealth follows the well-known boom–bust–recovery cycle of the broader crypto market, peaking near $1 T in late 2021 and tracking the TVL trajectory at consistently higher levels, before contracting sharply through 2022 as a sequence of market shocks unfolded.[3] From 2024 onwards the two series decouple, with portfolio wealth exhibiting an additional growth trend beyond TVL, potentially due to the rise of meme coins held outside DeFi protocols, and peaking at $1.47 T during 2025.

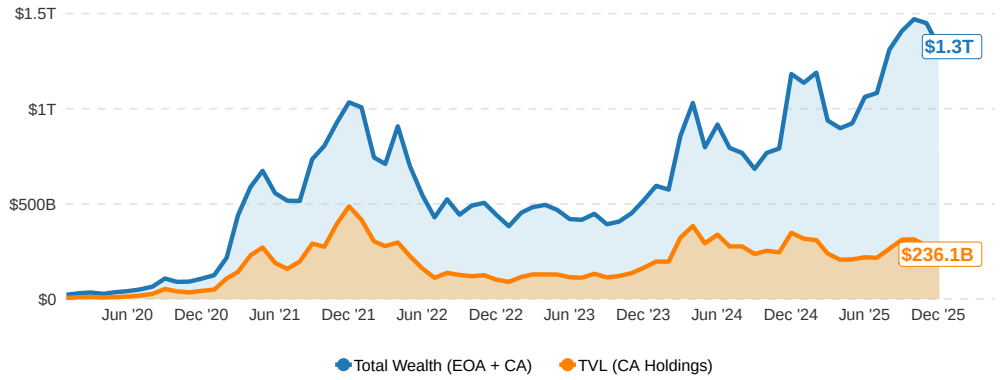

Figure 4 Total reconstructed portfolio value for all accounts over time, contrasted with TVL (total value locked) in DeFi protocols.

Disaggregating wealth by account type reveals a marked asymmetry: although CAs represent only 6.18 % of all asset-holding accounts, they hold on average 31.80% of total wealth, while EOAs account for the remaining 68.20%. The CA share is highest during the 2020–2021 DeFi expansion and declines toward end-2025 as wealth shifts to EOA accounts. This split is informative when comparing portfolio wealth with the TVL metric: TVL captures only value locked in DeFi protocols and therefore systematically underestimates total on-chain wealth by ignoring the large share held in EOA accounts. Although our reconstruction does not cover native ETH, the reconstructed account wealth still substantially exceeds TVL, indicating the magnitude of wealth not captured by the metric. Both series are also subject to double-counting of wrapped or derivative tokens [42].

4.2 Wealth Distribution and Concentration

We characterise the wealth distribution across accounts from both ends: first via shares of accounts by wealth bin and account type, then via the top 1% and 10% wealth shares as our primary measure of concentration, complemented by standard inequality indices.

As Figure 5 shows, the distribution of accounts across wealth bins differs systematically by account type but converges on the same headline: the majority of both CAs and EOAs

[3] Terra/LUNA (May 2022), Celsius and Three Arrows Capital (June–July 2022), FTX (November 2022).

consistently hold less than \$100, with CAs concentrated in the dust[4] bin ($\leq$ \$1, 51.63% share over the entire observation period) and EOAs in the \$1–\$100 bin (49.43%). The two types reach their low-wealth-share minima at different times, the CA minimum around mid-2021 ("DeFi Summer") and the EOA minimum around late 2021 at the broader market peak, jointly producing the two wealth peaks visible in Figure 4. Account growth trajectories, CA/EOA wealth shares over time, and the full bin-level distributions are reported in Appendix B.1.

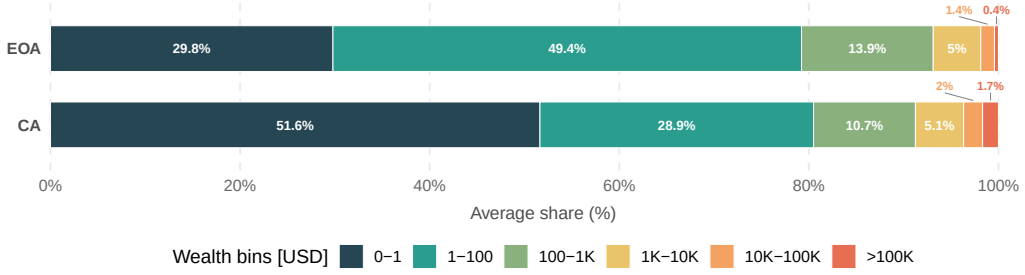

Figure 5 Share of EOA and CA accounts by wealth bucket (USD), computed across all 72 monthly snapshots.

Turning to the top of the distribution, we restrict the analysis to accounts whose portfolio value exceeds \$1, to mitigate the "dust" problem noted above.

We find that the top 1% of accounts consistently controls the vast majority of total portfolio value, increasing from 97.25% to 99.03% over 2020–2025. The top 10% likewise holds nearly the entire ecosystem value, always above 99.3%, while the number of accounts quadruples over the same period. These findings are corroborated by standard inequality measures (Gini, HHI), which remain uniformly high throughout the observation period. We treat these indices as supplementary rather than primary because address-level inequality measures are debated in pseudonymous blockchain settings, where one entity can control many addresses or one address can serve many entities [47, 10].

Narrowing to the top 0.1%, account type composition departs from the ecosystem average: CAs make up 18.95% of these addresses, far above their 6.18 % population share, with DeFi protocols, exchanges, and custodial wallets dominating the tagged categories. However, the aggregate CA value share within this slice (33.70%) is essentially unchanged from the ecosystem-wide 31.80%, meaning the same fraction of the pie is simply held by a proportionally larger group. The EOAs that reach the top 0.1% are correspondingly whales with highly concentrated individual wealth.

Yearly concentration statistics and account type breakdowns are provided in Appendix B, as well as the full Gini and HHI analysis and tag-based attribution details.

4.3 Portfolio Size

Having characterised wealth levels and concentration, we now examine portfolio size N, the number of distinct assets held, as a measure of diversification. Portfolio size serves as a proxy for diversification, as holding more assets reduces overall portfolio variance. We focus on accounts holding at least one asset at the snapshot.

[4] Millions of accounts can hold negligible residual balances after a trade, which can inflate the denominator of the distribution and mechanically increase measured concentration.

Figure 6 shows the distribution of portfolio sizes across wealth categories, averaged across all monthly snapshots. Overall, 83.35% of accounts hold only a single asset, 10.14% hold two, 4.26% hold three to four, and only 2.17% hold five or more distinct assets. Single-asset accounts dominate in every wealth category, but the share of multi-asset portfolios rises steadily with wealth, with the exception of the highest wealth category, where the trend reverses, likely reflecting the prevalence of CAs such as exchanges and liquidity pools that concentrate holdings in a small number of high-value tokens.

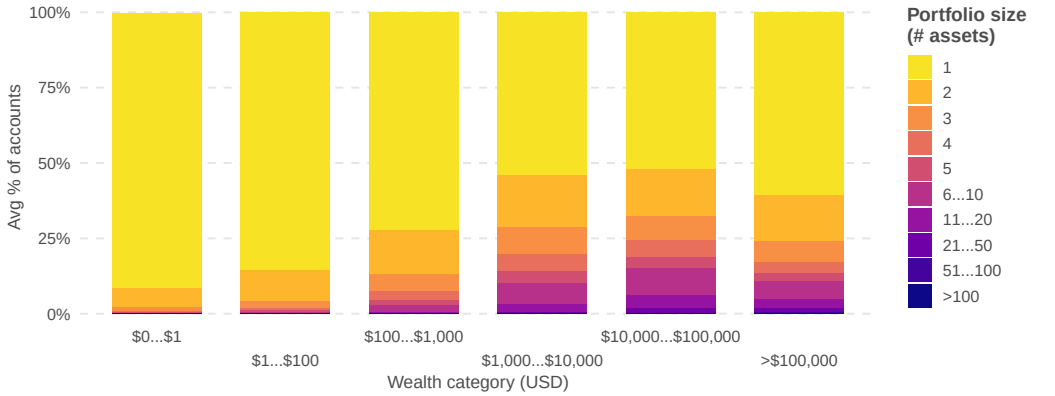

Figure 6 Portfolio size distribution: share of accounts holding a given number of distinct assets, by wealth category, averaged across all monthly snapshots.

Takeaways

Total on-chain wealth tracks the TVL trajectory at consistently higher levels, peaking at $1.47 T during 2025. CAs account for only 6.18 % of asset-holding addresses but hold 31.80% of aggregate wealth, so TVL systematically undercounts on-chain wealth by ignoring the EOA-held share. On-chain wealth is highly concentrated: the top 1% of accounts controls 99.03% in 2025, with the top 0.1% dominated by DeFi protocols, exchanges, and custodial wallets. While 83.35% of accounts hold only a single asset, the multi-asset share rises with wealth except at the very top. *Overall, on-chain portfolios are small along both dimensions: most accounts hold modest wealth (<$100) in only one or two tokens.*

5 Distance to the Efficient Frontier

This section tackles **Obj B** by measuring how far portfolios sit from the constrained efficient frontier and identifying what predicts their suboptimality. For every account holding at least two tokens, we solve the three mean-variance optimisation (MVO) problems by following strategies (MaxRet, MinVar, MaxSR) at 72 monthly snapshots (Jan 2020 – Dec 2025), yielding 239 634 756 (240 M) account-month optimisations. We provide supplementary analyses in Appendix C and computational details in Appendix E.

5.1 Measuring Portfolio Distance

We quantify the gap between actual ($w_{(0)}$) and optimal ($w_{(s)}$) allocations using the weight distance d defined in Def. (4), examining overall means, distributional patterns across the

three strategies, and temporal stability. We rescale the raw distance from $[0, 1]$ to $[0, 100]\%$ throughout, where 0% denotes coincidence with the optimum and 100% zero weight overlap.

Distance prevalences. Figure 7 groups the distances into buckets for each of the three optimisation strategies, reporting average shares over the entire observation period. The resulting distances reveal a clear ordering across strategies: on-chain portfolios are closest to the MinVar optimal solution (12.59% mean), followed by MaxRet (28.30%), and farthest from the MaxSR optimum (45.09%). However, these means mask important distributional differences: the three strategies show strikingly different profiles. MinVar and MaxRet are left-skewed, with the bulk of account-observations concentrated near the optimum (63.35% and 40.09% within 1%, respectively). In contrast, MaxSR is right-skewed, with 24.66% of observations in the $(60, 80]\%$ bin.

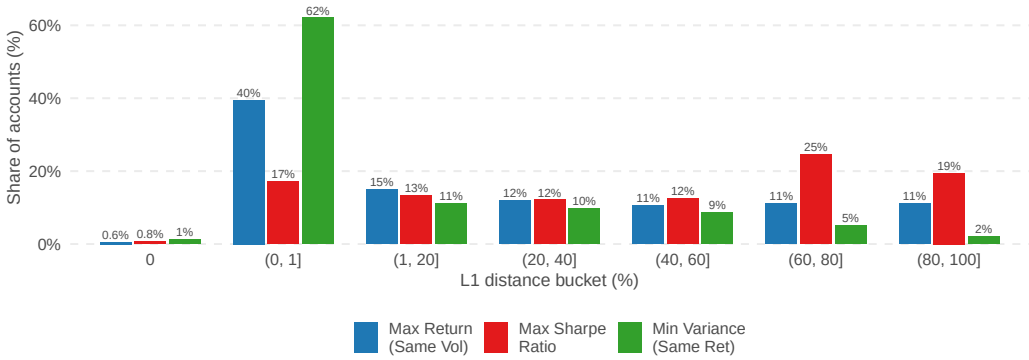

Figure 7 Distribution of portfolio distance to the constrained efficient frontier for the three optimisation strategies. Share of 240 M account-month optimisations across seven distance bins.

Temporal dynamics. Distance may vary with portfolio size and evolve with market conditions over time. To investigate, we stratify accounts at an illustrative threshold of $N = 5$ for portfolio size and track the mean distance over time in Figure 8, with one panel per optimisation strategy and standard-error bands around each trajectory. All three distances exhibit a slight upward trend over time.

Small portfolios ($N < 5$) have a consistently small average distance to MinVar (7.5 %), indicating near-structural alignment, while MaxSR shows the largest (41.5 %). Large portfolios ($N \geq 5$) show a sharp increase to 48.5 % for MinVar, and 70.2 % for MaxSR. A paired Wilcoxon signed-rank test across the 72 monthly snapshots confirms this small–large gap is significant with large effect sizes.

Notably, this near-optimal share is almost entirely due to two-asset portfolios, which have little room to deviate from the minimum-variance weights and therefore sit close to the optimum by construction: once we restrict to $N \geq 5$, the share within 1% of the optimum drops to near zero. Test statistics, Cohen's d values, per-size-bucket plots, threshold sensitivity, cross-strategy correlations, and robustness analysis are in Appendix C.

5.2 Determinants and Modeling of Distance

Having characterised how distance is distributed and evolves over time, we now ask what predicts suboptimality and whether it admits a parsimonious model. We first identify the dominant determinants, then formalise the distance–size relationship ($d \sim N$) with a parametric model, and finally benchmark actual allocations against naive strategies.

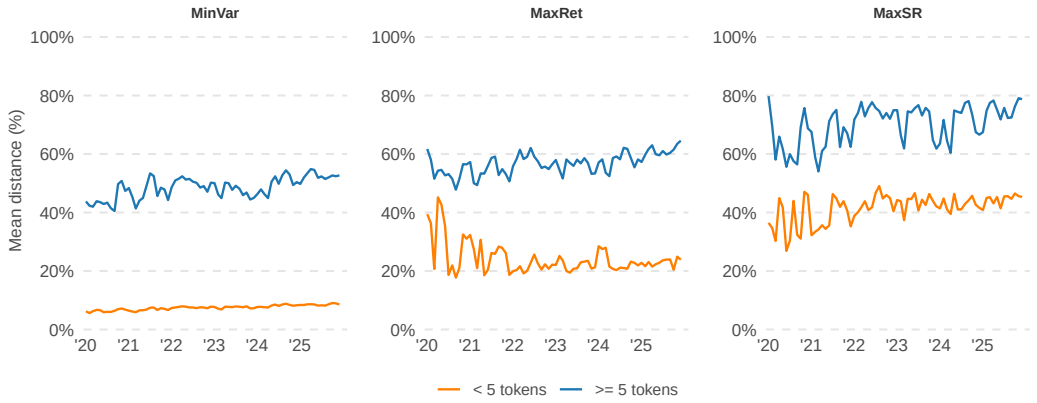

Figure 8 Mean distance over time by portfolio size. Trajectories for portfolios holding fewer than five assets versus five or more, with standard-error bands; one panel per optimisation strategy.

On a sampled subset, we regress distance (d) on three features: portfolio size (N), the log portfolio value ($\log V$), and the snapshot month (t). As shown in Figure 9, N is the dominant determinant of *mean* distance across all three strategies. Random Forest models show that N also dominates the *variance* of MinVar distance across individual accounts (95.5%), driven by a near-optimal cluster of almost exclusively two-asset portfolios, while for MaxRet and MaxSR portfolio value becomes the strongest individual-level predictor (45%).

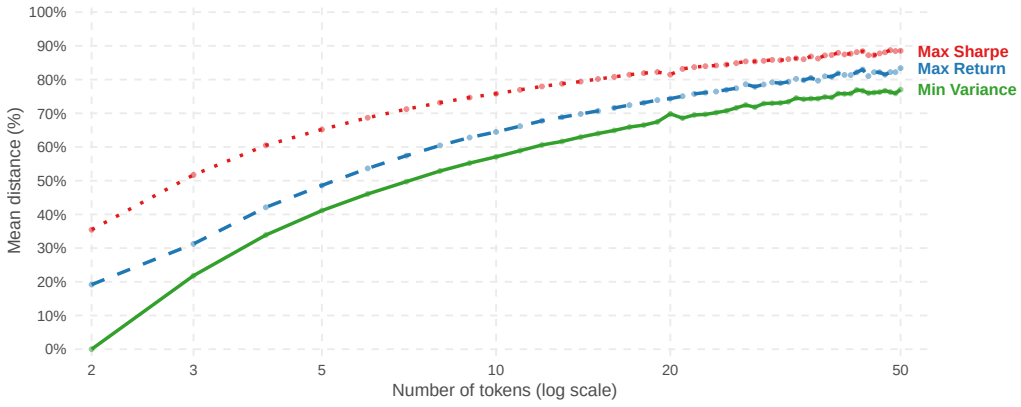

Figure 9 Mean distance versus portfolio size. Distance d as a function of size N on a log-scaled axis for all three optimisation strategies.

Parametric model of distance. As Figure 9 shows, the $d \sim N$ relationship is monotone and concave across all three strategies: distance rises steeply from $N = 2$ to $N = 5$ and plateaus thereafter, consistent with diminishing marginal deviation as portfolio size grows. This plateau also motivates the $N \geq 5$ stratification used in Section 5.1. We therefore formalise it with a parsimonious three-parameter power-decay model:

$$\hat{d}_s(N) \;=\; \delta_s^{\infty}\left(1 \;-\; \psi_s \cdot N^{-\gamma_s}\right), \tag{2}$$

where $\hat{d}_s$ is the fitted mean distance in percent, $N \geq 2$ the portfolio size, $\delta_s^{\infty} \leq 100\%$ the asymptotic ceiling, $\psi_s > 0$ a shape parameter setting the depth of the dip from the asymptote,

and $\gamma_s > 0$ the exponent governing how quickly distance saturates as N grows. All three fits achieve $R^2 \geq 0.99$ (Table 1). The asymptote shows that MaxRet reaches 100%, while MinVar and MaxSR plateau at 80.69% and 88.62%, implying persistent structural overlap with the minimum-variance and Sharpe-optimal allocations. Estimation details and fitted curves are in Appendix C.

Table 1 Power-decay distance model. Fitted parameters and goodness-of-fit for Equation (2), by optimisation strategy.

	Model Parameters			Model Quality	
Strategy s	ψ_s	γ_s	δ_s^∞ (%)	R^2	MAE (%)
MaxRet	1.1319	0.4803	100.00	0.9937	1.12
MinVar	1.7139	0.7774	80.69	0.9999	1.56
MaxSR	1.1211	0.9022	88.62	0.9999	1.34

Comparison with naive allocation benchmarks. To put these distance results into perspective, we compare the optimisation strategies against two naive allocation benchmarks that require no estimation of returns or covariances: an *equal-weight* portfolio ($w_i = 1/N$) and a *market-capitalisation-weight* portfolio. The mean distance to equal-weight is 37.7% and to market-capitalisation-weight 38.4%, both substantially larger than the 12.59% distance to MinVar. While both means are similar, the medians reveal a distributional difference: for equal-weight the median (40.0%) exceeds the mean (37.7%), whereas for market-capitalisation-weighted (31.2% median, 38.4% mean) the reverse holds. The typical account is thus noticeably closer to market-capitalisation-weight than to equal-weight portfolios, which is intuitive: concentrated holdings in one or two large-cap tokens mechanically resemble a market-cap-weighted portfolio, even without deliberately targeting the weighting.

This benchmark comparison examines the structural effect of MPT optimisation, revealing that the three strategies shift accounts in qualitatively distinct ways. MaxRet and MinVar both produce a negative mean distance-delta against equal-weight, meaning optimisation on average moves accounts *closer* to an equal-weight allocation; MinVar additionally leaves 63.5% of accounts unchanged, consistent with many portfolios already lying near the constrained minimum-variance solution. MaxSR, by contrast, pushes the majority of accounts farther from both equal-weight and cap-weight (full decomposition in Appendix C).

Takeaways

Across the three strategies, portfolios are closest to the minimum-variance optimum (MinVar, mean 12.59%) and farthest from the Sharpe-optimal frontier (MaxSR, 45.09%). The large near-MinVar share (distance $\leq 1\%$: 63.35%) is structural: two-asset portfolios mechanically approximate the minimum-variance solution, and the share collapses for larger portfolios ($N \geq 5$). Portfolio size N is the dominant determinant of the mean distance to the constrained efficient frontier across all three strategies. Larger portfolios exhibit substantial distance to all three frontiers, indicating considerable room for rebalancing. Relative to naive benchmarks, actual portfolios are nearer to market-cap-weighting than to equal-weighting, reflecting concentrated holdings in large-cap tokens. MaxRet and MinVar optimisation moves accounts *closer* to equal-weight, whereas MaxSR pushes them farther from both naive benchmarks. *Overall, on-chain portfolios sit closest to the minimum-variance frontier, but this proximity is largely explained by the dominant two-asset portfolio size.*

6 Realised Portfolio Performance

Finally, we address **Obj C** by quantifying how deviations from the constrained efficient frontier translate into realised performance relative to the market. For each account snapshot, we compute the $T^{(+)}$-day realised return under the actual, optimised, and naive allocations.

6.1 Forward Returns and Risk-Adjusted Returns

We set $T^{(+)} = 20$ days, keeping adjacent forward returns approximately non-overlapping [22] and yielding a 3:1 $T^{(-)}/T^{(+)}$ ratio standard in rolling mean–variance evaluation [14, 27]. The actual allocation $w_{(0)}$ serves as baseline against the optimised MPT strategies (MaxRet, MinVar, MaxSR) and naive allocations (equal-weight, market-cap-weight). Table 2 reports, for the forward-realised return $r_{t'}^{(a,t)}$, the median and the hit rate, which is the fraction of accounts beating the baseline. Alongside, the risk-adjusted return α is documented with the median and the fraction of accounts with positive results for each strategy. Throughout, we take the median within each monthly snapshot and weight snapshots equally, so that cohort-size growth does not dominate the statistics.

Table 2 Realised returns and risk-adjusted performance by strategy. Median forward return over the $T^{(+)}$-day horizon, hit rate versus the baseline allocation, and CAPM α.

Strategy	Median Return (%)	Hit Rate (%)	Median α (%)	Positive α (%)
Baseline (actual account)	-1.87	—	-1.91	38.9
MaxRet	-2.33	48.5	-2.31	37.8
MinVar	-2.00	51.0	-1.85	38.5
MaxSR	-1.92	49.8	-1.92	38.0
Equal-weight ($1/N$)	-2.82	50.4	-2.78	37.9
Market-cap-weight	-0.24	53.3	-0.30	39.7
Market benchmark	+1.89	—	—	—

While the market benchmark r_m reaches 1.89%, the baseline realised-return is -1.87% for the actual allocation. Market-cap-weighting achieves the highest hit rate (53.3%) and the least negative median return (-0.24%). MinVar and equal-weight also exceed a 50% hit rate (51.0% and 50.4%, respectively), with MaxSR and MaxRet only marginally below.

The risk-adjusted return is negative for every allocation. Market-cap-weighting has the least negative median α (-0.30%) and the highest positive-α fraction (39.7%), consistent with the CAPM prediction that a cap-weighted allocation is hardest to beat. Among the MPT strategies, only MinVar is marginally less negative ($\alpha = -1.85\%$ vs. baseline -1.91%).

Aggregating per snapshot rather than averaging across all accounts in the observation period gives the same picture: across the six strategies, the median across snapshots of the share of accounts beating the wETH–WBTC market index varies only from 36.05% to 42.98%, and the share with positive α from 36.07% to 40.23%.

6.2 Predictors of Performance

To identify which account characteristics predict performance, we train Random Forest models on six features: entry month, portfolio value, portfolio size N, portfolio β, and binary indicators for holding ETH or BTC wrappers.

Market entry month dominates both forward and risk-adjusted returns, accounting for 70–79% of the explained variance for realised returns ($R^2 = 0.56$–0.65) and 55–62% for α.

Figure 10 confirms this pattern under repeated monthly optimisation: the cumulative excess return over the market benchmark trends progressively more negative for every strategy across the sample period. Portfolio β is the second strongest signal, modest for realised returns (14–22%) and more pronounced for α (25–31%), while portfolio value, size N, and asset composition contribute marginally.

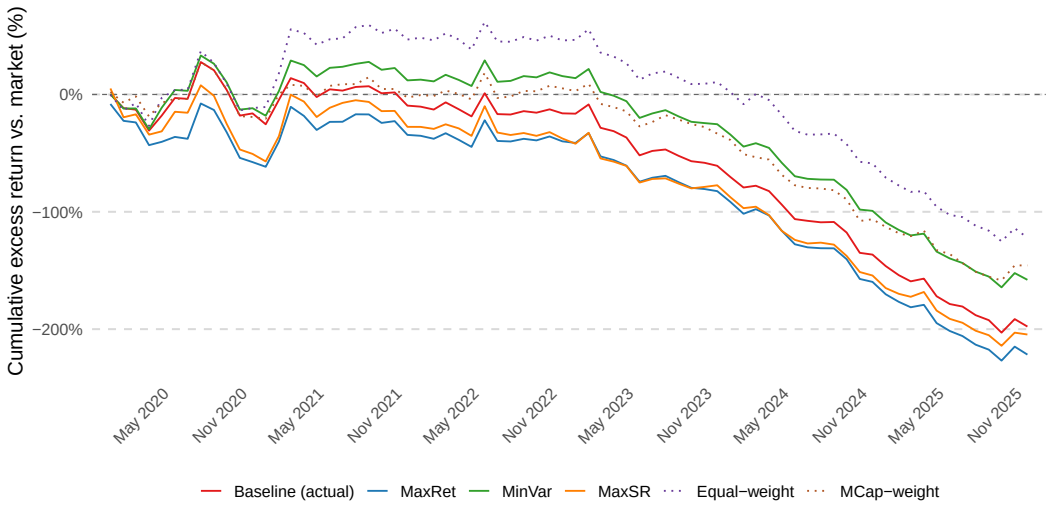

Figure 10 Cumulative excess return over the market benchmark. Running sum of monthly median return differentials by strategy, with allocations re-optimised independently at each snapshot.

We furthermore provide Random Forest, SHAP, ℓ_1-distance–α, and per-snapshot robustness checks in Appendices D.1, D.3, and D.4.

Takeaways

The market benchmark outperforms all optimisation strategies, including MPT, in median return. The dominant predictor of performance is market entry timing: the age of an account explains 70–79% of return variance and 55–62% of α variance, far exceeding the contribution of allocation choice, portfolio size, or asset composition. *Overall, when an account enters the market matters far more than how it allocates.*

7 Discussion

7.1 Key Findings

We present the first large-scale empirical study comparing actual on-chain portfolios to their theoretically optimal allocations under MPT. We propose a portfolio reconstruction method, apply it to Ethereum, and compute portfolios for 116 M accounts across 4562 ERC-20 tokens.

First, we analyse on-chain portfolios and find that wealth is highly skewed: the top 1% of accounts control 99.03% of token wealth in 2025, and most accounts (83.35%) hold a single asset. These results are consistent with prior work on narrow sets of assets (e.g., governance tokens [2, 37, 39]) and extend that evidence to a substantially larger and more heterogeneous set of tokens. High concentration thus persists across token categories. Relative to traditional finance, single-asset dominance on Ethereum is markedly more pronounced than the median US retail holding of three to four stocks [20]. We also find little evidence of the naive $1/N$

heuristic [3], under which allocations are divided equally across available options. On-chain allocation behaviour therefore departs from a pattern that is common in traditional retirement and investment settings.

Second, we measure how far portfolios sit from the constrained efficient frontier. Small portfolios lie closest to the minimum-variance optimum, but this proximity is structural rather than deliberate: portfolios near the optimum are predominantly two-token portfolios that mechanically approximate the minimum-variance allocation. Apparent optimality erodes as portfolios grow; for five or more tokens, mean deviations from the frontier reach 48.5%–70.2%, depending on the optimisation strategy. This extends prior evidence that low-dimensional allocations leave little room for optimisation to add value [14]. We further formalise the relationship between portfolio size and mean distance to the theoretical optima through a parsimonious three-parameter power-decay model.

Third, we quantify how deviations from the efficient frontier translate into realised performance relative to the market. Passive market-capitalisation weighting outperforms portfolios constructed under any of the three investigated MPT optimisation strategies. The dominant predictor of differences in realised returns is *when* an account enters the market, not *how* it allocates: entry month alone explains 70–79% of the variance in returns and 55–62% of the variance in risk-adjusted returns, far exceeding the contribution of portfolio size or allocation choice. Moreover, even under repeated monthly optimisation, all strategies accumulate increasingly negative excess returns relative to the market benchmark; this underperformance is therefore not confined to a single market regime. These results stand in contrast to established findings from traditional finance, which attribute roughly 90% of the time-series variation in portfolio returns to asset allocation policy [7, 25]. A plausible explanation is offered by Liu et al. [30], who show that a small number of common factors (in particular a cryptocurrency market factor) capture the variation in returns across cryptoassets, leaving limited scope for diversification benefits across tokens.

Taken together, these findings challenge the descriptive adequacy of MPT in the cryptoasset domain. Decentralised markets are natural empirical testbeds for MPT; actual allocations, however, bear little resemblance to mean-variance optimal portfolios. Pervasive under-diversification suggests that risk-return optimisation is not a salient feature at the wallet level. The dominance of entry timing over allocation choice further indicates that returns are governed primarily by a common market factor [30], a regime in which the diversification benefits MPT promises are structurally limited. Mean-variance optimisation thus appears neither descriptive of observed behaviour nor prescriptively useful in this setting, even if MPT remains a normative benchmark.

Our findings have practical implications. Our empirical analysis suggests that rebalancing toward MPT optima does not reliably improve risk-adjusted performance, while passive market-cap weighting is observed to offer a more robust default. These findings could be of use to whoever wishes to import the mean-variance framework to the blockchain setting (e.g., *retail investors or wealth managers*): doing so wholesale risks projecting an illusion of optimisation onto a return regime in which a dominant market factor leaves little for diversification gains. For *platform and protocol designers*, the persistence of single-asset portfolios among smaller accounts suggests that frictions (e.g., gas fees, limited attention) shape observed behaviour more than the celebrated Modern Portfolio Theory (MPT); interventions targeting transaction costs or portfolio-construction interfaces may therefore have greater practical impact than tools encouraging frontier-based optimisation. For *regulators and supervisors*, the extreme wealth concentration we document raises concerns about systemic exposure that allocation-focused frameworks do not address.

7.2 Limitations and Future Work

Our analysis is subject to several limitations, each of which points to directions for future work. First, we cover only ERC-20 token portfolios on Ethereum, restricted to tokens listed on Coingecko. This subset excludes scams, abandoned projects, and illiquid assets, but is biased toward established tokens. Native ETH enters the dataset indirectly via its wrapped representation (WETH) [18], the form of ETH usable in DeFi protocols. Assets held on custodial exchanges are not observable on-chain and are therefore excluded. This scope is analogous to prior work on traditional portfolios, which likewise focuses on specific markets and instruments (e.g., [20, 24]). Extending the reconstruction to other EVM-compatible blockchains and to native-token flows can broaden coverage along this axis.

Second, we treat each Ethereum address as an independent portfolio. Linking addresses to common entities remains an open problem in blockchain analytics [47, 10, 29]; off-chain heuristics such as address clustering cannot guarantee full accuracy [36, 45]. Entity-level portfolios may therefore differ in size and composition from the address-level portfolios we observe, and yield a complementary view at the actor level.

Third, our study evaluates portfolios against MPT benchmarks, but does not analyse the underlying decision-making process (e.g., trading motives, herding, disposition effects). While most accounts underperform the market benchmark, characterising the minority of outperformers could reveal whether systematic strategies or informational advantages exist on-chain. The publicly available reconstruction pipeline and dataset (upon publication) provide a foundation for such analyses.

8 Conclusion

We reconstructed token portfolios for 116 M Ethereum accounts across 4562 ERC-20 tokens over the full chain history (2015–2025) and evaluated them against MPT benchmarks. Both the actual on-chain portfolios and their MPT-optimal counterparts underperform the market benchmark in realised returns. On-chain account portfolios are pervasively under-diversified, sitting close to the constrained efficient frontier by construction. Realised returns are more closely tied to *when* accounts enter the market rather than *how* they allocate. MPT thus appears neither descriptive nor prescriptively useful in this setting, even as it retains its value as a normative benchmark. The public reconstruction pipeline and dataset (upon publication) enable entity-level analyses, cross-chain extension, and outperformer characterisation.

References

1 Brad M. Barber and Terrance Odean. Trading is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors. *The Journal of Finance*, 55(2):773–806, 2000. doi:10.1111/0022-1082.00226.

2 Tom Barbereau, Reilly Smethurst, Orestis Papageorgiou, Alexander Rieger, and Gilbert Fridgen. DeFi, not so decentralized: The measured distribution of voting rights. In *Proceedings of the 55th Hawaii International Conference on System Sciences (HICSS)*, pages 6043–6052, 2022. doi:10.24251/HICSS.2022.734.

3 Shlomo Benartzi and Richard H. Thaler. Naive diversification strategies in defined contribution saving plans. *American Economic Review*, 91(1):79–98, February 2001. doi:10.1257/aer.91.1.79.

4 Board of Governors of the Federal Reserve System (US). 3-Month Treasury Bill Secondary Market Rate, Discount Basis [TB3MS], 2025. Retrieved from FRED, Federal Reserve Bank of St. Louis. URL: https://fred.stlouisfed.org/series/TB3MS.

5 Nicola Borri, Yukun Liu, Aleh Tsyvinski, and Xi Wu. Cryptocurrency as an Investable Asset Class: Coming of Age, October 2025. arXiv:2510.14435.

6 Marie Brière, Kim Oosterlinck, and Ariane Szafarz. Virtual currency, tangible return: Portfolio diversification with Bitcoin. *Journal of Asset Management*, 16(6):365–373, 2015. doi:10.1057/jam.2015.5.

7 Gary P. Brinson, L. Randolph Hood, and Gilbert L. Beebower. Determinants of portfolio performance. *Financial Analysts Journal*, 42(4):39–44, 1986. doi:10.2469/faj.v42.n4.39.

8 Laurent E. Calvet, John Y. Campbell, and Paolo Sodini. Down or out: Assessing the welfare costs of household investment mistakes. *Journal of Political Economy*, 115(5):707–747, October 2007. doi:10.1086/524204.

9 John Y. Campbell. Household finance. *The Journal of Finance*, 61(4):1553–1604, 2006. doi:10.1111/j.1540-6261.2006.00883.x.

10 Nir Chemaya, Aviv Yaish, Shahar Yacouel, Dahlia Malkhi, and Lin William Cong. Quantifying inequality in blockchain networks, September 2025. URL: https://papers.ssrn.com/abstract=5540258.

11 Coingecko. Coingecko: Cryptocurrency prices and market capitalization. https://www.coingecko.com. Accessed: 2026-01-19.

12 Jeffrey Dean and Sanjay Ghemawat. MapReduce: Simplified data processing on large clusters. In *Proceedings of the 6th Symposium on Operating Systems Design and Implementation (OSDI)*, pages 137–150. USENIX Association, 2004.

13 DefiLlama. DefiLlama – DeFi dashboard. https://defillama.com/, 2025. Accessed: 2025-12-01.

14 Victor DeMiguel, Lorenzo Garlappi, and Raman Uppal. Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *The Review of Financial Studies*, 22(5):1915–1953, May 2009. doi:10.1093/rfs/hhm075.

15 Tatiana Didier, Roberto Rigobon, and Sergio L. Schmukler. Unexploited Gains From International Diversification: Patterns Of Portfolio Holdings Around The World. *The Review of Economics and Statistics*, 95(5):1562–1583, oct 2013. doi:10.1162/REST_a_00351.

16 Edwin J. Elton, Martin J. Gruber, Stephen J. Brown, and William N. Goetzmann. *Modern Portfolio Theory and Investment Analysis*. Wiley, Hoboken, USA, 9 edition, 2014. URL: https://openlibrary.org/books/OL25441475M/Modern_portfolio_theory_and_investment_analysis.

17 Erigon. Erigon (formerly turbo-geth). https://github.com/erigontech/erigon. Accessed: 2026-01-19.

18 ethereum.org. What is wrapped ether (WETH), 2026. Accessed: 2026-04-22. URL: https://ethereum.org/en/wrapped-eth/.

19 Kenneth R. French and James M. Poterba. Investor diversification and international equity markets. *American Economic Review*, 81(2):222–226, 1991. Papers and Proceedings.

20 William N. Goetzmann and Alok Kumar. Equity Portfolio Diversification. *Review of Finance*, 12(3):433–463, 2008. doi:10.1093/rof/rfn005.

21 GraphSense. GraphSense REST api, 2024. Accessed: 2025. URL: https://github.com/graphsense/graphsense-REST.

22 Lars Peter Hansen and Robert J. Hodrick. Forward exchange rates as optimal predictors of future spot rates: An econometric analysis. *Journal of Political Economy*, 88(5):829–853, 1980. doi:10.1086/260910.

23 Oliver Hellum, Theis Ingerslev Jensen, Bryan T. Kelly, and Semyon Malamud. Complex modern portfolio theory. SSRN Working Paper 6443319, March 2026. URL: https://ssrn.com/abstract=6443319, doi:10.2139/ssrn.6443319.

24 Gur Huberman. Familiarity breeds investment. *The Review of Financial Studies*, 14(3):659–680, 2001. doi:10.1093/rfs/14.3.659.

25 Roger G. Ibbotson and Paul D. Kaplan. Does asset allocation policy explain 40, 90, or 100 percent of performance? *Financial Analysts Journal*, 56(1):26–33, 2000. doi:10.2469/faj.v56.n1.2327.

26 Iknaio Cryptoasset Analytics. Iknaio cryptoasset address tags, 2026. Accessed: 2025. URL: https://www.iknaio.com/.

27 Raymond Kan and Guofu Zhou. Optimal portfolio choice with parameter uncertainty. *Journal of Financial and Quantitative Analysis*, 42(3):621–656, 2007. doi:10.1017/S0022109000004129.

28 Olivier Ledoit and Michael Wolf. Honey, i shrunk the sample covariance matrix. *The Journal of Portfolio Management*, 30(4):110–119, July 2004. doi:10.3905/jpm.2004.412318.

29 Qiangqiang Liu, Qian Huang, Frank Fan, Haishan Wu, and Xueyan Tang. Detecting sybil addresses in blockchain airdrops: A subgraph-based feature propagation and fusion approach. In *2025 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*. IEEE, 2025. arXiv:2505.09313, doi:10.1109/ICBC64466.2025.11185061.

30 Yukun Liu, Aleh Tsyvinski, and Xi Wu. Common Risk Factors in Cryptocurrency. *The Journal of Finance*, 77(2):1133–1177, 2022. doi:10.1111/jofi.13119.

31 Scott M. Lundberg, Gabriel Erion, Hugh Chen, Alex DeGrave, Jordan M. Prutkin, Bala Nair, Ronit Katz, Jonathan Himmelfarb, Nisha Bansal, and Su-In Lee. From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, 2(1):56–67, 2020. doi:10.1038/s42256-019-0138-9.

32 Scott M. Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, volume 30, pages 4765–4774, 2017.

33 Harry Markowitz. Portfolio Selection. *The Journal of Finance*, 7(1):77–91, 1952. doi:10.1111/j.1540-6261.1952.tb01525.x.

34 Harry M. Markowitz. Foundations of Portfolio Theory. *The Journal of Finance*, 46(2):469–477, 1991. doi:10.1111/j.1540-6261.1991.tb02669.x.

35 Massimo Massa and Andrei Simonov. Hedging, Familiarity and Portfolio Choice. *Review of Financial Studies*, 19(2):633–685, 2006. doi:10.1093/rfs/hhj013.

36 Sarah Meiklejohn, Marjori Pomarole, Grant Jordan, Kirill Levchenko, Damon McCoy, Geoffrey M. Voelker, and Stefan Savage. A fistful of bitcoins: Characterizing payments among men with no names. In *Proceedings of the 2013 Internet Measurement Conference (IMC)*, pages 127–140. ACM, 2013. doi:10.1145/2504730.2504747.

37 Matthias Nadler and Fabian Schär. Decentralized finance, centralized ownership? An iterative mapping process to measure protocol token distribution. *Journal of Blockchain Research*, 1(1):29–36, March 2022. doi:10.4310/JBR.2022.v1.n1.a3.

38 Terrance Odean. Are Investors Reluctant to Realize Their Losses? *The Journal of Finance*, 53(5):1775–1798, 1998. doi:10.1111/0022-1082.00072.

39 Christina Ovezik, Dimitris Karakostas, Mary Milad, Aggelos Kiayias, and Daniel W. Woods. SoK: Measuring blockchain decentralization. In *Applied Cryptography and Network Security (ACNS 2025)*, volume 15825 of *LNCS*. Springer, 2025. arXiv:2501.18279, doi:10.1007/978-3-031-95761-1_7.

40 Alla Petukhina, Simon Trimborn, Wolfgang Karl Härdle, and Hermann Elendner. Investing with cryptocurrencies: Evaluating their potential for portfolio allocation strategies. *Quantitative Finance*, 21(11):1825–1853, 2021. doi:10.1080/14697688.2021.1880023.

41 Emmanouil Platanakis and Andrew Urquhart. Portfolio management with cryptocurrencies: The role of estimation risk. *Economics Letters*, 177:76–80, 2019. doi:10.1016/j.econlet.2019.01.019.

42 Pietro Saggese, Michael Fröwis, Stefan Kitzler, Bernhard Haslhofer, and Raphael Auer. Towards verifiability of total value locked (TVL) in decentralized finance, 2025. arXiv:2505.14565, doi:10.48550/arXiv.2505.14565.

43 William F. Sharpe. Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk. *The Journal of Finance*, 19(3):425–442, 1964. doi:10.1111/j.1540-6261.1964.tb02865.x.

44 Simon Trimborn, Mingyang Li, and Wolfgang Karl Härdle. Investing with Cryptocurrencies: A liquidity constrained investment approach. *Journal of Financial Econometrics*, 18(2):280–306, 2020. doi:10.1093/jjfinec/nbz016.

45 Friedhelm Victor. Address clustering heuristics for ethereum. In *Financial Cryptography and Data Security: 24th International Conference, FC 2020 , Kota Kinabalu, Malaysia, February 10–14, 2020 Revised Selected Papers*, pages 617–633, Berlin, Heidelberg, 2020. Springer-Verlag. doi:10.1007/978-3-030-51280-4_33.

46 Friedhelm Victor and Bianca Katharina Lüders. Measuring ethereum-based ERC20 token networks. In *Financial Cryptography and Data Security*, volume 11598 of *LNCS*, pages 113–129, Berlin, Heidelberg, 2019. Springer-Verlag. doi:10.1007/978-3-030-32101-7_8.

47 Aviv Yaish, Nir Chemaya, Dahlia Malkhi, and Lin William Cong. Inequality in the age of pseudonymity. In *Proceedings of the Fortieth AAAI Conference on Artificial Intelligence (AAAI-26)*, AAAI'26/IAAI'26/EAAI'26, pages 17293–17301. AAAI Press, 2026. arXiv:2508.04668.

A Glossary

Following is a list of the terms, notations, and acronyms used in the paper.

A.1 Glossary

covariance matrix	Matrix Σ whose entry (i, j) is the covariance of returns between tokens i and j; the diagonal contains per-token variances. Estimated empirically over a lookback window and used as input to MVO.
cryptoasset	Any cryptocurrency or digital token that can be transferred and traded on a blockchain.
efficient frontier	The set of portfolios that achieve the highest expected return for a given level of volatility. Throughout the paper, the empirically relevant quantity is the *constrained* efficient frontier: we optimise over each account's already-held tokens rather than over all available tokens, and native ETH enters only through its wrapped form (WETH).
Gini coefficient	A measure of statistical dispersion ranging from 0 (perfect equality) to 1 (maximal inequality), used here to quantify wealth inequality among token holders.
naive allocation	A baseline allocation that requires no estimation of returns or covariances. The $1/N$ (equal-weight) variant assigns weight $1/N$ to each of the N held tokens; the market-cap variant weights each token by its market capitalisation. Used here as benchmarks against MPT-optimal allocations.
Sharpe ratio	The ratio of a portfolio's expected excess return to its volatility, measuring risk-adjusted performance.
token	A digital asset issued on a blockchain platform, typically following a standard such as ERC-20 on Ethereum.

A.2 Acronyms

CA	contract account
CAGR	compound annual growth rate
CAPM	capital asset pricing model
DeFi	decentralised finance
EOA	externally owned account
ERC-20	Ethereum request for comments 20
HHI	Herfindahl-Hirschman index
MaxRet	max return at target volatility
MaxSR	max Sharpe ratio
MinVar	min variance at target return
MPT	modern portfolio theory
MVO	mean-variance optimisation
TVL	total value locked
WBTC	wrapped BTC
wETH	wrapped ETH

B Supplementary Material: Portfolio Characteristics

This appendix provides supplementary analysis for the portfolio characteristics discussed in Section 4. It contains additional detail on account evolution and growth trajectories (Section B.1), the decomposition of top holders by account type (Section B.2), address-level attribution of the wealthiest accounts (Section B.3), and standard inequality measures confirming the wealth concentration documented in the main text (Section B.4).

B.1 Accounts and Wealth

This subsection details account growth trajectories, the distribution of wealth across bins and account types, and the concentration of value among top holders, supplementing the descriptive statistics in Section 4.

B.1.1 Account evolution.

We analyse the evolution of accounts that have acquired at least one asset of our token selection. We further distinguish between EOAs and CAs in the line plot of Figure 11.

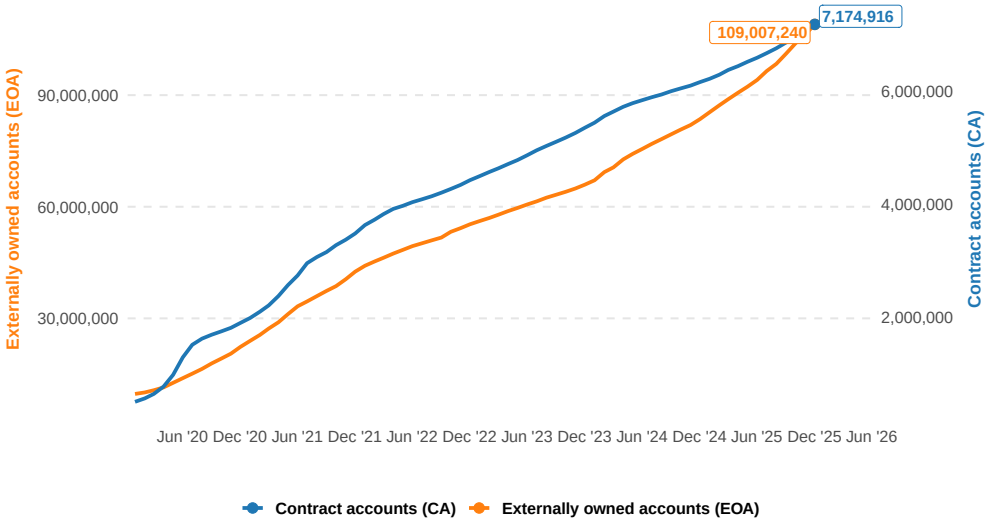

Figure 11 Number of accounts: the temporal evolution of accounts that have acquired at least one asset.

The two account types exhibit substantially different growth trajectories. By the end of 2025, EOAs reach 109 007 240 addresses against 7 174 916 CAs, so that contract accounts represent only 6.18 % of all asset-holding accounts. The EOA population grows at a compound annual growth rate (CAGR) of 50.53 % over 2020–2025, adding roughly 16 785 778 addresses per year. CAs grow at a comparable CAGR of 56.93 %, corresponding to approximately 1 128 397 new contracts per year, reflecting the steady deployment of DeFi protocols and token infrastructure.

B.1.2 Wealth distribution of accounts.

We classify accounts into wealth bins

$$(0, 1],\ (1, 100],\ (100, 1\text{K}],\ \ldots,\ (100\text{K}, \infty) \quad [\text{USD}],$$

restricting the statistics to active addresses (non-zero wealth); this binning scheme reflects the heavy-tailed nature of on-chain wealth distributions. Table 5 reports the time-averaged share of accounts in each bin separately for CAs and EOAs, while Figures 12 and 13 show the temporal evolution for the low- and high-wealth bins respectively.

Table 5 Wealth bin distribution: time-averaged share of accounts in each wealth bin across the observation window (non-zero USD value accounts only). Shares are percentages.

Type	0–1	1–100	100–1K	1K–10K	10K–100K	>100K
CA	51.63	28.86	10.75	5.08	2.00	1.67
EOA	29.80	49.43	13.89	5.02	1.44	0.42

Across both account types, the majority of accounts consistently hold wealth below \$100, but the within-type composition differs markedly: for CAs the lowest-wealth bin ($\leq$ \$1) holds the largest share (51.63%), whereas for EOAs the \$1–\$100 bin is the largest (49.43%). The share of low-wealth accounts decreases from 2020 to 2021 and rises thereafter, while the share of high-wealth accounts peaks around mid-2021, consistent with the market run-up and contraction described above. The two account types reach their respective minima at different times: the CA minimum around mid-2021 ("DeFi Summer"), whereas the EOA minimum around late 2021 aligns with the broader market peak, together producing the two wealth peaks visible in Figure 4.

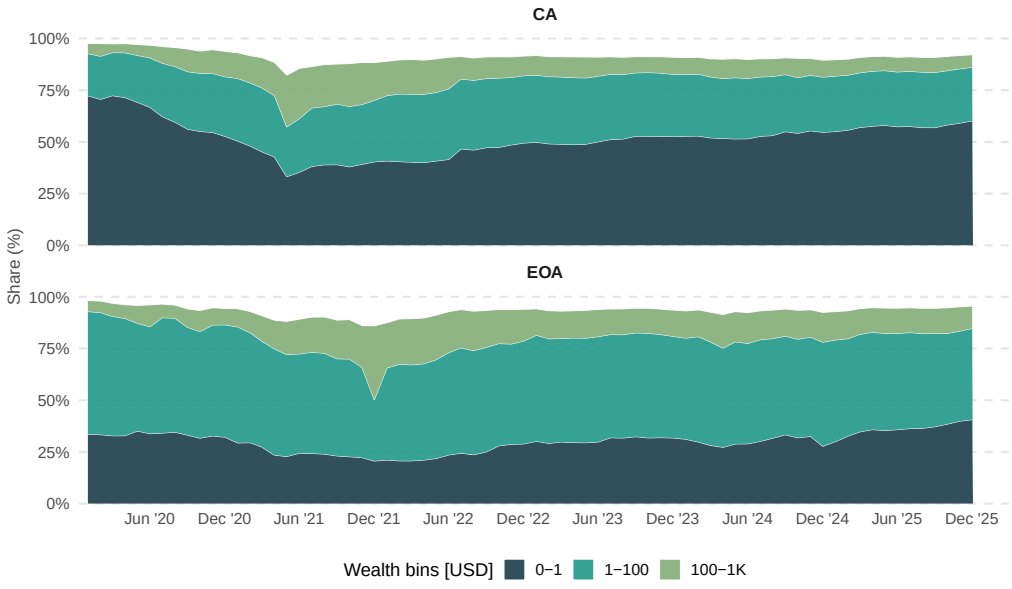

Figure 12 Distribution of accounts across *low-wealth* bins over time (share of accounts with non-zero wealth). The top subplot shows CA, the bottom subplot EOA.

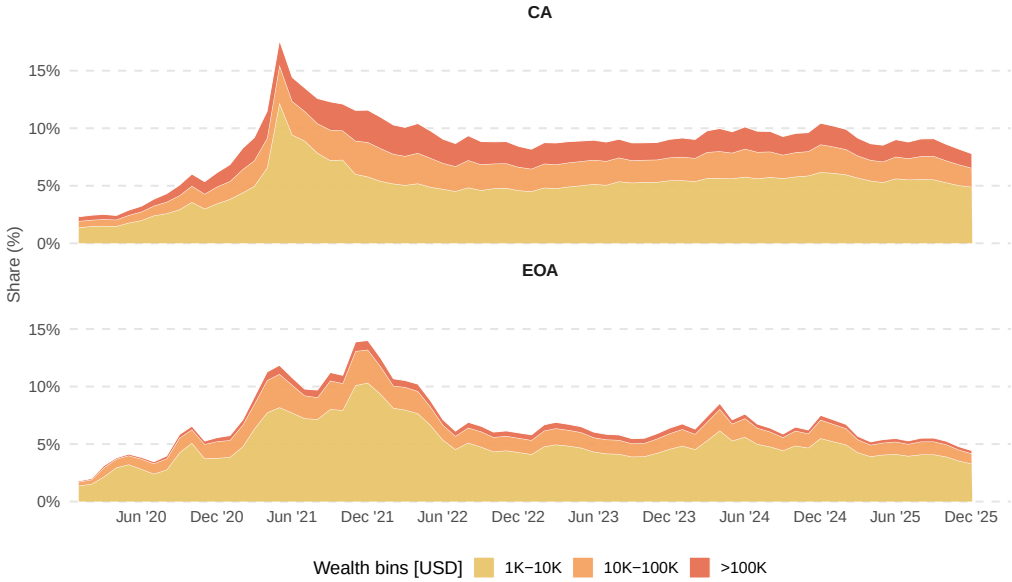

Figure 13 High-wealth bin distribution over time: share of accounts across high-wealth bins, with CA (top) and EOA (bottom). The peak in high-wealth CA accounts around mid-2021 coincides with DeFi Summer, when protocols experienced rapid TVL growth through liquidity mining.

B.1.3 Wealth of top accounts.

Table 6 summarises the yearly concentration statistics for all accounts with non-zero wealth, extending the aggregate figures reported in Section 4.2.

Table 6 Top-holder concentration by year: yearly averages of top-holder wealth shares for all accounts. Counts are in millions; shares in percent. The Top 1% (min–max) column captures within-year variation across monthly snapshots.

Year	Avg accounts (M)	Top 1% (avg)	Top 1% (min–max)	Top 5% (avg)	Top 10% (avg)
2020	5.003	97.25	96.40–98.80	99.37	99.75
2021	10.025	97.66	96.73–98.12	99.40	99.73
2022	14.277	98.16	97.91–98.36	99.49	99.73
2023	14.747	97.79	97.66–97.91	99.51	99.78
2024	18.646	98.39	97.97–98.68	99.57	99.80
2025	21.058	99.03	98.69–99.28	99.74	99.87

B.2 Account Type Breakdown of Top Holders

To understand whether the extreme wealth concentration documented in Section 4.2 is driven by protocol-controlled accounts or individual holders, we decompose the top 1% and top 0.1% by account type.

Within the top 1%, contract accounts represent only 6.98% of accounts by count but hold 32.26% of the top-1% wealth on average. EOAs account for the remaining 93.02% of accounts and 67.74% of value. The CA share is broadly consistent with the ecosystem-wide CA wealth share of 31.80% reported in Section 4.1, indicating that the top-1% is not disproportionately

composed of smart contracts relative to the overall population.

Narrowing the focus to the top 0.1% sharpens the picture. CAs now constitute 18.95% of addresses, nearly three times their share in the top 1%, while their value share remains stable at 33.70%. This indicates that the very largest holders include a substantially higher proportion of smart contracts, likely corresponding to DeFi protocol treasuries, exchange hot wallets, and token vesting contracts, while the majority of wealth even at this level is still held by individual EOA addresses.

B.3 Address Attribution of Top-0.1% Holders

To further characterise the entities behind the wealthiest addresses, we match the top 0.1% of accounts at each monthly snapshot against address tags provided by iknaio [26], queried via the GraphSense REST API [21]. The dataset assigns category labels and semantic concepts to known Ethereum addresses.

On average, 23.25% of top-0.1% addresses carry at least one tag (range: 10.70–35.14% across snapshots), while the remaining 76.20% are unlabelled. The limited coverage reflects the pseudonymous nature of Ethereum: the majority of high-value addresses have no publicly attributed identity. Among the $\approx 24\%$ with attribution, two complementary views emerge.

By category. Of the addresses carrying a GraphSense *category* label (9.6% of all top-0.1% block–address pairs), the largest group is *user* (33.83% of categorised), followed by *defi_dex_pair* (19.32%), *service* (15.66%), *exchange* (7.97%), and *defi_token* (7.40%). Together, exchange- and DeFi-related categories account for over 40% of categorised addresses.

By concept. The broader *concept* taxonomy covers 21.4% of top-0.1% pairs. The dominant concept is *defi* (63.07% of concept-tagged addresses), followed by *defi_custody* (31.73%), *exchange* (17.09%), and *defi_dex* (13.88%). This confirms that the top of the wealth distribution is heavily populated by DeFi protocol contracts (liquidity pools, custody vaults, token treasuries) and centralised exchange wallets, consistent with the disproportionate CA value share reported in Section B.2.

B.4 Wealth Inequality and Concentration

Given that the large majority of accounts hold less than $100 on average, we further investigate the wealth distribution to quantify the degree of concentration. We use the Herfindahl-Hirschman Index (HHI) to measure the concentration of wealth among accounts and Gini coefficients to analyse the concentration and inequality of token portfolios.

Before turning to these indices, we briefly characterise the token-level holder base, which provides context for interpreting concentration measures. Over time, the typical token holder base shrinks: the average number of holders per token decreases from 27 088 (median: 3312) in 2020 to 12 429 (median: 1746) in 2025. This decline is driven by a compositional shift. Figure 14 summarises the annual distribution of tokens grouped by their holder count. The share of tokens with more than 10 000 holders drops from 28.3% to 15.6%, while smaller tokens (100–1 000 holders) grow from 17.7% to 27.5%. Mid-sized tokens (1 000–10 000 holders) remain the dominant group throughout, ranging from 39.7% to 50.4%.

To complement the top-holder analysis in Section 4.2, we compute two standard inequality measures for each token at every monthly snapshot and report their temporal evolution. We restrict the analysis to tokens with more than 100 holders at the time of the snapshot to focus on mature assets.

The *Gini coefficient* ranges from 0 (all accounts hold equal wealth) to 1 (one account holds all wealth). The *Herfindahl–Hirschman Index* (HHI) is calculated as $\mathrm{HHI} = \sum_{a=1}^{N} s_a^2$,

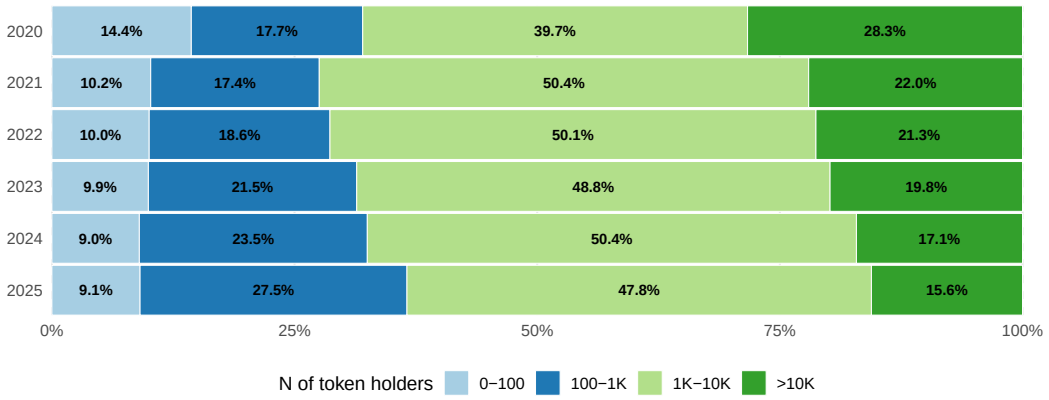

Figure 14 Token holder distribution: share of Ethereum request for comments 20 (ERC-20) tokens grouped by number of holders per token (2020–2025).

where s_a is the wealth share of account a and N is the total number of accounts; higher values indicate greater concentration.

Figure 15 displays the average and median of both indices across all qualifying tokens over time. Gini coefficient coefficients remain uniformly close to 1 throughout the observation period, with the cross-token average ranging from 0.954 to 0.978 and the median from 0.975 to 0.990. HHI values indicate moderate to high concentration, with the average ranging from 0.221 to 0.297 and the median from 0.139 to 0.220. Both measures are stable over time, confirming that the extreme wealth concentration documented by the top-holder shares in the main text is a persistent structural feature of the ERC-20 token ecosystem.

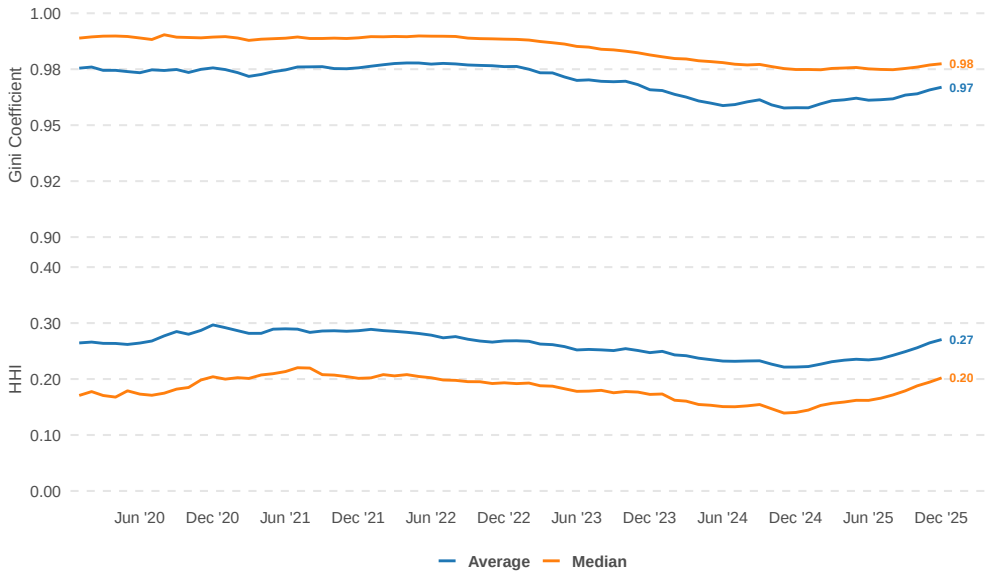

Figure 15 Inequality measures over time: average and median Gini coefficient coefficient (top) and Herfindahl-Hirschman index (HHI) (bottom) across all tokens with more than 100 holders, computed at each monthly snapshot (2020–2025).

C Supplementary Material: Distance to the Efficient Frontier

This appendix supplements the distance analysis in Section 5. It covers the formal definition of the weight ℓ_1 distance d_s and the mean-reversion regularisation applied to return estimates (Section C.1), the calibration and robustness of the risk-free rate r_f (Section C.2), the determinants of d_s and cross-strategy robustness checks (Section C.3), the relationship between d_s and portfolio size with its fitted power-decay model (Section C.4), the comparison with naive allocation benchmarks (Section C.5), and the optimisation volume and temporal evolution of d_s (Section C.6).

C.1 Weight ℓ_1 Distance and Mean-Reversion Regularisation

This subsection provides the formal definition of the ℓ_1 weight distance used throughout the main analysis (Section 5), together with the mean-reversion regularisation applied to return estimates in the optimisation pipeline.

Formal definition.

Let

$$w_{(0)} = \big(w_{(0),1}, \ldots, w_{(0),N}\big), \qquad w_{(s)} = \big(w_{(s),1}, \ldots, w_{(s),N}\big),$$

with $\sum_{i=1}^{N} w_{(0),i} = \sum_{i=1}^{N} w_{(s),i} = 1$ and $w_i \geq 0$ (long-only). We use ℓ_1 because it has a direct *rebalancing interpretation*: for fully-invested portfolios, $\frac{1}{2}\|w_{(0)} - w_{(s)}\|_1$ equals the minimum one-way turnover (fraction of value traded) required to move from the current to the optimised weights. We therefore define the distance with the $\frac{1}{2}$ factor included (see Definition 4).

Using $|u - v| = u + v - 2\min\{u, v\}$ for $u, v \geq 0$ and the fact that both portfolios are fully invested, we obtain

$$d_s = \tfrac{1}{2}\left\|w_{(0)} - w_{(s)}\right\|_1 \;=\; 1 - \sum_{i=1}^{N} \min\left\{\, w_{(0),i},\, w_{(s),i} \,\right\}. \tag{3}$$

For long-only portfolios, the distance is bounded by

$$0 \;\leq\; d_s \;\leq\; 1, \tag{4}$$

where 0 indicates identical allocations and values close to 1 indicate nearly disjoint allocations (mass invested in different assets).

Mean-reversion regularisation.

Sample mean returns over the $T^{(-)}$-day lookback are noisy in cryptoasset markets, where extreme short-term moves are common and tend to be partially corrected over subsequent windows. We exploit this mean-reversion property as a regularisation mechanism: by shrinking each asset's sample mean μ_i toward a cross-sectional prior, we reduce sensitivity to transient price spikes and short-lived momentum effects. Formally, for each asset i we replace the raw sample mean return μ_i (Definition 1) by the shrunk estimate

$$\mu_i^{\text{shrunk}} = \lambda\,\bar{\mu} \;+\; (1 - \lambda)\,\mu_i, \tag{5}$$

where $\bar{\mu} = \frac{1}{N}\sum_{j=1}^{N}\mu_j$ is the cross-asset mean and $\lambda \in [0,1]$ is the shrinkage strength. In our configuration we set $\lambda = 0.5$, corresponding to a moderate pull toward the prior. The resulting portfolio expected return is then computed as usual by

$$\mu^{(a,t)} = \sum_{i=1}^{N} w_i^{(a,t)} \mu_i^{\text{shrunk}},$$

which reduces overconfident forecasts while still preserving meaningful cross-sectional differences in recent performance. In practice, this produces smoother optimised weights and more stable allocations across snapshots.

C.2 Calibration and Robustness of the Risk-Free Rate

The MaxSR (Definition 3) is known as the *tangency portfolio*, constructed by drawing a line from the risk-free rate r_f tangent to the constrained efficient frontier. It is therefore the only place in our optimisation pipeline where r_f enters. The MinVar and MaxRet optima depend only on the covariance structure and the moments of the actual portfolio, and the CAPM alphas are evaluated at $r_f = 0$ to keep the cross-sectional comparison invariant under a uniform shift in the riskless return. We adopt $r_f = 5\%$ per annum throughout the MaxSR optimisation on three independent grounds.

Reference to the literature.

The closest published precedent on portfolio optimisation with crypto assets [41] proxies the risk-free rate with the U.S. 3-month Treasury bill yield. Their convention has carried over to subsequent crypto-portfolio work, and our setting matches it directly: in the 2023–2024 window, which contains the bulk of our account-time observations, the FRED series `TB3MS` averaged $\approx 5\%$ [4].

Empirical match to on-chain stablecoin yields.

Account holders in our sample interact natively with on-chain markets rather than U.S. Treasuries, so we cross-check the T-bill anchor against its closest on-chain analogue: the variable-rate supply APY for USDC and USDT on the major non-custodial money markets active on Ethereum during the sample window. We aggregate the supply APYs reported by DefiLlama's yields API across Aave v2 and v3, Compound v2 and v3, Morpho Blue, and Spark, restricting to pools deployed on Ethereum with TVL above one million USD and APY in $[0, 50]\%$ to exclude inactive or anomalous pools, and average the resulting series within each calendar month from January 2021 onwards. Table 7 reports the unweighted monthly statistics across all months in which lending data are available.

Table 7 On-chain stablecoin supply APY across the major non-custodial money markets on Ethereum (Aave v2/v3, Compound v2/v3, Morpho Blue, Spark), aggregated from DefiLlama's yields API. Unweighted monthly statistics, in %.

Asset	Avg APY	Median monthly APY	Months observed
USDC	5.51	4.77	44
USDT	4.47	3.48	52
USDC and USDT pooled	4.94	3.88	52

The pooled mean of 4.94% essentially matches the T-bill anchor of 5%. We retain the T-bill convention of [41] for comparability across the literature, but note that the on-chain alternative produces essentially the same number.

Robustness to the value of r_f.

The asymmetry between $r_f = 5\%$ in the MaxSR optimisation and $r_f = 0$ in the CAPM alpha evaluation is a deliberate modelling choice that we now show is empirically harmless on the MaxSR side. We re-run the full optimisation pipeline with $r_f = 0$ on a stratified random sample of 5,000 accounts per monthly snapshot (360,000 accounts across all 72 snapshots) and compare the resulting MaxSR weights and distance metrics, account-by-account, against the production results obtained with $r_f = 5\%$. Table 8 summarises the per-account differences.

Table 8 Robustness of the MaxSR results to the value of r_f. Per-account quantities computed across 360,000 matched accounts, normalised to the 0–100 scale used throughout the paper.

Quantity	Median	Mean	95th pct (abs)
distance between MaxSR weights ($r_f = 0$ vs. $r_f = 5\%$)	0	4.5	37.1
Δd_{MSR} ($w_{(0)}$ vs. MaxSR)	0	-0.4	13.2
Δd (MaxSR vs. equal-weight)	0	-0.7	2.4
Δd (MaxSR vs. market-cap-weight)	0	-2.6	18.5

The median account sees no change in any reported quantity. The mean changes in the three distance metrics are well under three percentage points in absolute value. At the 95th percentile only the MaxSR weights themselves shift meaningfully (37.1%), reflecting that the tangency portfolio slides along the same constrained efficient frontier when r_f changes. The distance d_{MSR} from the actual allocation $w_{(0)}$ to that frontier barely moves. Crucially, the change in the headline distance d_{MSR} has a mean of -0.4 pp and a 95th-percentile absolute change of 13.2 pp on a 0–100 scale.

C.3 Determinants and Robustness of Distance

We investigate which observable account characteristics determine distance using the feature set `num_tokens`, `log_value_usd`, and `month` on a 10-million-row sample.

Regression and feature importance.

OLS. Linear models explain little variance for MaxRet ($R^2 = 0.064$) and MaxSR ($R^2 = 0.080$), but substantially more for MinVar ($R^2 = 0.358$). In all three, `num_tokens` carries the largest coefficient ($p < 0.001$). An extended specification with polynomial and interaction terms raises R^2 to 0.104, 0.132, and 0.632 respectively. The dominant additional predictor is log(`num_tokens`), confirming the logarithmic relationship between portfolio complexity and distance visible in Figure 9.

Random Forest. A 300-tree model on a held-out set achieves $R^2 = 0.376$ (MAE = 19.55 pp) for MaxRet, $R^2 = 0.676$ (MAE = 5.35 pp) for MinVar, and $R^2 = 0.412$ (MAE = 19.35 pp) for MaxSR. Feature importances, shown in Figure 16, diverge sharply. For MinVar, portfolio size dominates with 95.5% of total importance. For the other two, importance is spread more evenly across `log_value_usd` (∼45%), `num_tokens` (27–31%), and `month` (25–28%). This confirms the qualitative split visible in Figure 9: MinVar distance is almost

entirely determined by portfolio size, whereas MaxRet and MaxSR distances additionally depend on wealth and market regime.

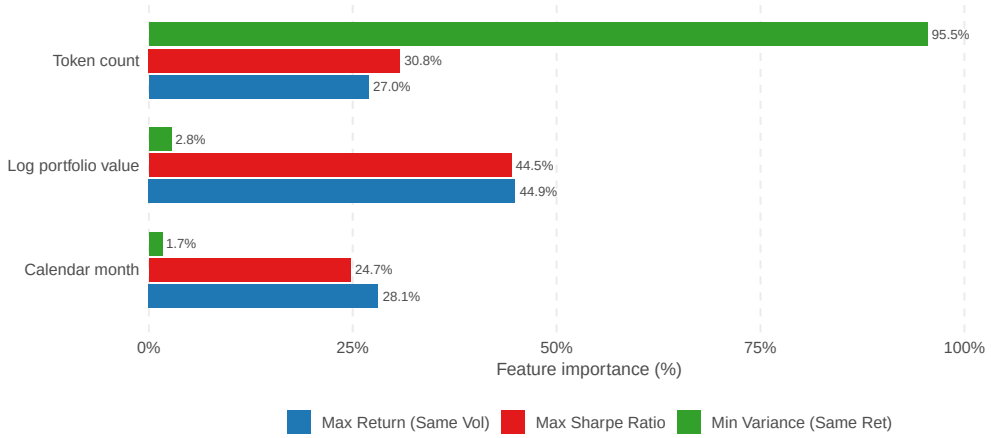

Figure 16 Random Forest feature importances: permutation-based importances for predicting distance under each optimisation strategy. Portfolio size accounts for 95.5% of importance under MinVar but only 27–31% under the other two objectives.

Spearman correlations on a 50M-row sample confirm the non-linearity: $\rho = 0.811$ between `num_tokens` and MinVar distance, far exceeding the Pearson r of 0.576. For MaxRet and MaxSR the Spearman ρ is 0.451 and 0.399, positive but less dominant.

Quantile regression reveals that at $\tau = 0.25$, only `num_tokens` is significant for MaxRet and MinVar (low-distance accounts hold exactly two tokens regardless of value), while at $\tau = 0.75$ the `log_value_usd` coefficient also becomes significant, indicating that wealth matters primarily in the high-distance tail.

Near-optimal accounts.

We examine accounts whose distance falls within 1% of the optimum. Across all three objectives, near-optimal accounts are overwhelmingly simple: 92–95% hold exactly two tokens, and cumulative three-or-fewer exceeds 99% in every case. Median portfolio value is modest ($27–$33). The key distinction is the *size* of the near-optimal pool: 151.8M observations (63.4%) are near-optimal under MinVar, 96.1M (40.1%) under MaxRet, and only 42.8M (17.9%) under MaxSR, confirming that minimum-variance optimality is structurally easy to approximate, while Sharpe optimality is difficult even for the simplest allocations.

Per-strategy distance by portfolio size.

Figure 17 disaggregates the relationship between portfolio size and distance for each optimisation strategy individually, showing the median (line), mean (dots), interquartile range (shaded band), and ±1 standard-deviation envelope. The standard-deviation bands reveal substantial cross-sectional heterogeneity at every portfolio size, particularly for MaxSR, where the upper band exceeds 80% even for portfolios with $N > 10$. For MinVar, the standard deviation narrows rapidly with portfolio size, confirming that larger portfolios converge tightly toward the minimum-variance optimum.

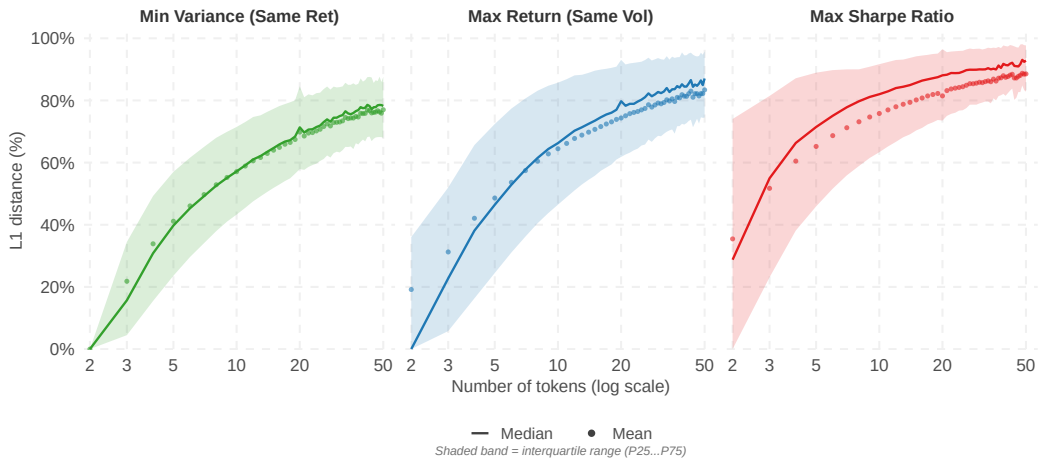

Figure 17 Per-strategy distance by portfolio size: median distance (line), mean (dots), interquartile range (shaded band), and ±1 standard-deviation envelope as a function of portfolio size N (log-scaled axis). Based on a 5M-row sample.

Cross-strategy correlation.

The three distance measures are only moderately correlated with one another: $r = 0.452$ between MaxRet and MinVar, $r = 0.360$ between MaxRet and MaxSR, and $r = 0.385$ between MinVar and MaxSR. Being optimally allocated for one objective thus provides limited information about optimality for another, so portfolio "quality" in DeFi is not a unidimensional concept.

Joint optimality is exceedingly rare. Only 3099 account-observations (0.001%) achieve exact zero distance under all three objectives simultaneously. Pairwise exact-zero overlap is also thin: 0.26% for MaxRet ∩ MinVar, 0.05% for MaxRet ∩ MaxSR, and 0.03% for MinVar ∩ MaxSR.

Robustness to portfolio size.

The preceding results are heavily shaped by the two-asset majority. To test whether the findings generalise to genuinely diversified portfolios, we repeat the full analysis on the subset of account-observations with $N \geq 5$. The near-optimal population essentially vanishes: the share within 1% of the optimum drops from 63.35% to 0.49% for MinVar, from 40.09% to 0.77% for MaxRet, and from 17.87% to 0.16% for MaxSR, confirming that the high near-optimality rates in the full dataset were an artefact of portfolio simplicity rather than genuine optimisation. The distribution also shifts markedly rightward: under MaxSR, 50.6% of $N \geq 5$ accounts fall in the $(80, 100]$% bucket, compared to 19.49% in the full data.

Predictive models lose most of their explanatory power. Random Forest R^2 for MinVar drops from 0.676 to 0.173, and the Spearman ρ between `num_tokens` and MinVar distance falls from 0.811 to 0.365, indicating that the previously dominant portfolio-size effect was largely capturing the gap between two-asset and multi-asset accounts. Feature importances reshuffle accordingly: for MaxSR, `month` becomes the leading predictor (42.7%), suggesting that shifting market regimes make the Sharpe target a moving goalpost for diversified portfolios. Cross-strategy correlations also increase (most notably, r between MaxRet and MaxSR rises from 0.360 to 0.596), implying that among diversified accounts, suboptimality across

objectives converges. Overall, these results reinforce the paper's central finding: observable portfolio characteristics explain suboptimality primarily through the coarse channel of portfolio size, while *which* tokens are held (unobserved in our feature set) is the dominant predictive feature for any non-trivial allocation.

C.4 Distance by Portfolio Size

Figure 18 disaggregates the binary split from Figure 8 into eight portfolio-size buckets. Across all three strategies, mean distance rises steeply as portfolio size increases from $N = 2$ to $N = 5$ and then plateaus for larger portfolios. For MinVar, $N = 2$ accounts have a mean distance near zero, while $N = 5$ accounts already approach the level observed for the largest portfolios ($N \geq 21$). The pattern is qualitatively similar for MaxRet and MaxSR, though the absolute distances are higher and the plateau sets in at a slightly higher level. This saturation motivates the choice of $N = 5$ as the stratification threshold in the main analysis: below $N = 5$, distance is a mechanical artefact of portfolio simplicity, while above $N = 5$, additional tokens contribute little further deviation.

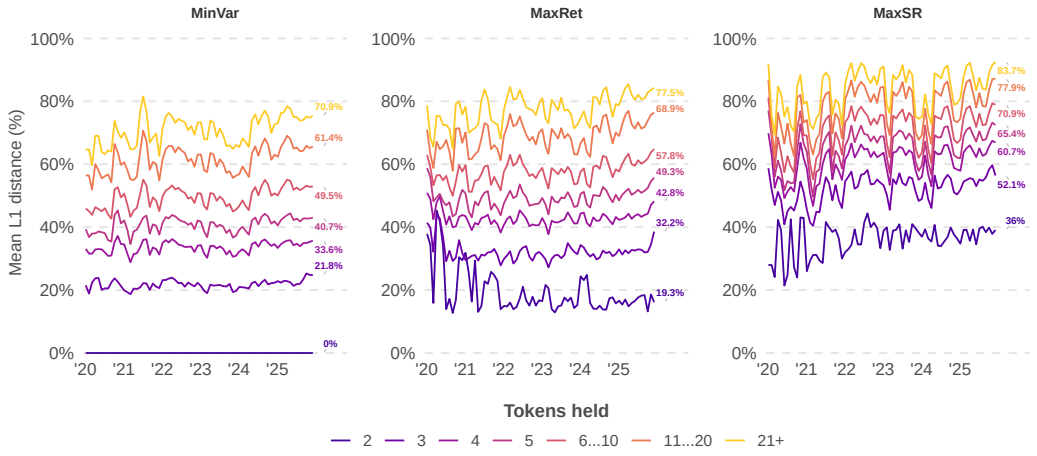

Figure 18 Mean distance over time by granular portfolio-size bucket and optimisation strategy. Legend entries report the time-series average for each bucket. Distance rises steeply from $N = 2$ to $N = 5$ and plateaus thereafter.

Threshold sensitivity.

To verify that the binary split is robust to the choice of cutoff, we repeat the paired Wilcoxon signed-rank test from Section 5.1 for thresholds $N \in \{3, 4, 5, 6, 7\}$. Table 9 reports the results. All fifteen tests reject the null at $p < 1.7 \times 10^{-13}$ (the minimum attainable for 72 paired blocks), confirming that the small-versus-large contrast is not an artefact of the specific cutoff. At $N = 5$, the threshold used in the main text, Cohen's d is 4.93 for MaxRet, 13.77 for MinVar, and 5.48 for MaxSR, the last two reflecting the near-zero distances achievable by small portfolios under minimum-variance optimisation. Cohen's d increases monotonically with N because raising the threshold purifies the "below" group (dominated by one- and two-token accounts that are structurally close to optimal), thereby widening the gap. The marginal gain in d is largest between $N = 3$ and $N = 5$ and diminishes thereafter, consistent with the saturation visible in Figure 18.

Table 9 Paired Wilcoxon signed-rank tests for the distance gap between portfolios with fewer than N tokens and those with N or more, at varying thresholds. All p-values equal 1.66×10^{-13} (Wilcoxon floor for 72 blocks).

Strategy	N	Mean$<N$ (%)	Mean$\geq N$ (%)	Δ (pp)	Cohen's d
MaxRet	3	19.3	42.4	23.0	3.39
	4	22.3	51.4	29.1	4.54
	5	24.2	56.8	32.6	4.93
	6	25.3	60.7	35.4	5.30
	7	26.1	63.6	37.5	5.59
MinVar	3	0.0	32.8	32.8	12.62
	4	4.9	42.7	37.8	13.84
	5	7.5	48.5	41.1	13.77
	6	9.0	52.6	43.6	14.18
	7	10.0	55.7	45.8	14.50
MaxSR	3	36.0	59.7	23.7	4.46
	4	39.7	66.5	26.9	5.22
	5	41.5	70.2	28.7	5.48
	6	42.6	72.8	30.2	5.79
	7	43.3	74.6	31.3	6.02

Power-decay fit.

We model the relationship between portfolio size and mean distance using a three-parameter power-decay specification $\hat{d}_s(n) = \delta_s^{\infty}\,(1 - \psi_s \cdot n^{-\gamma_s})$, estimated by weighted nonlinear least squares (Levenberg–Marquardt) on 49 portfolio-size bins ($n \in [2, 50]$, minimum 30 observations per bin), with the constraint $\delta_s^{\infty} \leq 100\%$. Each bin is weighted by $\sqrt{N_n}$, where N_n is the number of account-observations at that portfolio size, so that well-populated bins (e.g. n=2 with 38M observations) dominate the fit over sparse tail bins. The model is monotonically increasing and bounded: it equals $\delta_s^{\infty}(1 - \psi_s \cdot 2^{-\gamma_s})$ at the minimum portfolio size and asymptotes to δ_s^{∞} as $n \to \infty$. The asymptote δ_s^{∞} need not equal 100%: MaxRet hits the upper bound ($\delta^{\infty} = 100.00\%$), confirming that sufficiently diversified wallets can share zero weight overlap with the return-maximising portfolio.

In contrast, MinVar and MaxSR plateau at $\delta^{\infty} = 80.69\%$ and $\delta^{\infty} = 88.62\%$ respectively, indicating that even infinitely diversified wallets retain some structural overlap with the minimum-variance and Sharpe-optimal allocations. The exponent γ_s controls the saturation speed. For MaxSR, $\gamma = 0.9022$ is the largest, meaning that distance rises steeply with just a few additional tokens and largely plateaus beyond $n \approx 10$. For MaxRet, $\gamma = 0.4803$ is the smallest, indicating a more gradual climb that continues to increase noticeably up to $n \approx 20$. MinVar sits in between ($\gamma = 0.7774$), consistent with the intermediate position of minimum-variance distance throughout our analysis.

Figure 19 overlays the fitted curves on the empirical means. The fits track the data closely across the entire range, with the largest deviations appearing in the sparse tail beyond 30 tokens where individual bin estimates are noisier.

C.5 Naive Benchmark Comparison

Table 10 reports the distance between the actual allocation $w_{(0)}$ and two naive benchmarks. Both naive strategies are substantially farther from $w_{(0)}$ than the closest MPT strategy

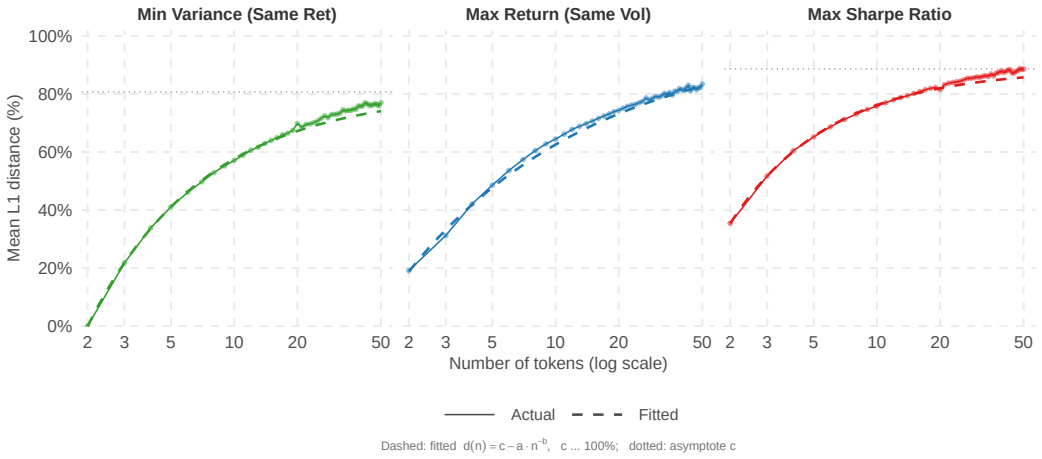

Figure 19 Fitted versus empirical distance: empirical mean distance (solid) versus fitted power-decay model $\hat{d}_s(n) = \delta_s^\infty\,(1 - \psi_s \cdot n^{-\gamma_s})$ (dashed) for each strategy. Dotted lines mark the asymptote δ_s^∞ where $\delta_s^\infty < 100\%$.

(MinVar), confirming that on-chain accounts do not simply replicate either equal-weight or market-cap-weight allocations. The mean–median relationship differs across benchmarks: for equal-weight the median exceeds the mean, indicating a left-skewed distribution with a concentration of accounts near the 40.0% mark, whereas for market-cap-weight the mean exceeds the median, reflecting a right tail of accounts that deviate substantially from cap-weighting.

Table 10 Distance to naive benchmarks: distance between the actual allocation $w_{(0)}$ and naive allocation benchmarks, expressed as percentages. The closest MPT strategy (MinVar) is included for reference.

Benchmark	Mean d_s (%)	Median d_s (%)
Equal-weight ($1/N$)	37.7	40.0
Market-cap-weight	38.4	31.2
MinVar (reference)	*12.59*	*0.0*

Table 11 and Figure 20 decompose the structural effect of each MPT strategy on distance to the naive benchmarks. Three patterns stand out. First, MinVar is the most conservative optimiser: it leaves the majority of accounts unchanged against both equal-weight (63.5%) and cap-weight (63.0%), and when it does move accounts, it predominantly moves them closer to equal-weight (closer-to-farther ratio $\approx$ 3:1). Second, MaxRet shows a split personality: it moves accounts closer to equal-weight (mean $\Delta d = -0.042$) but farther from cap-weight (mean $\Delta d = +0.130$), indicating that return-maximising portfolios tend to diversify away from concentrated large-cap holdings. Third, MaxSR pushes the majority of accounts farther from both benchmarks, with over half moving away from equal-weight (50.8%) and cap-weight (52.0%), reflecting the highly specific asset weightings required to maximise the Sharpe ratio.

Table 11 Directional effect of MPT optimisation on naive-benchmark distance: $\Delta d < 0$ indicates the optimised portfolio moved closer; $\Delta d > 0$ farther. Accounts with $|\Delta d| < 0.001$ are classified as unchanged.

Strategy	Naive Benchmark	Mean Δd	Closer (%)	Farther (%)	Unchanged (%)
MaxRet	Equal-weight	−0.042	32.9	20.1	46.9
	MCap-weight	+0.130	22.7	37.1	39.9
MinVar	Equal-weight	−0.074	27.4	9.1	63.5
	MCap-weight	−0.008	18.8	17.9	63.0
MaxSR	Equal-weight	+0.173	10.2	50.8	39.1
	MCap-weight	+0.205	29.8	52.0	17.9

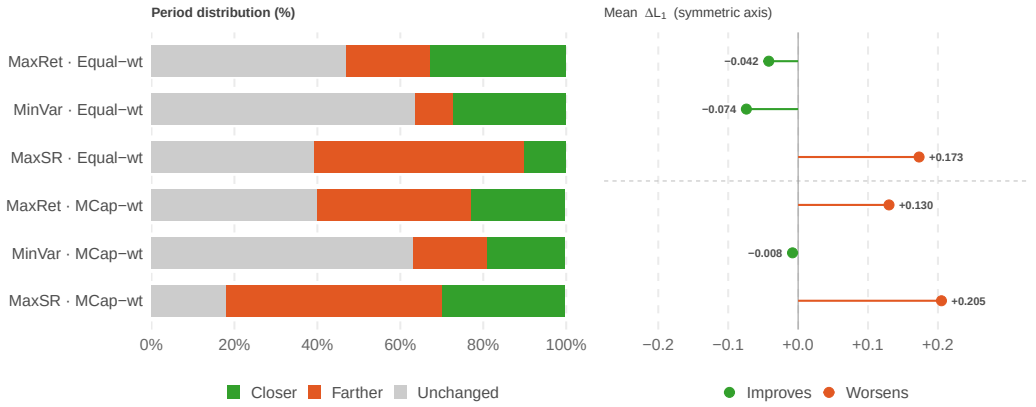

Figure 20 MPT optimisation versus naive benchmarks: left, proportion of accounts moved closer, farther, or left unchanged; right, mean Δd (negative = toward the benchmark; positive = away).

C.6 Optimisation Volume and Temporal Evolution

This subsection reports the optimisation volume and temporal distance statistics underlying the main analysis in Section 5.

Aggregate temporal evolution of distance.

Figure 21 reports the aggregate mean and median distance, together with ±1 standard-deviation bands, for each optimisation strategy over the full sample period. All three strategies exhibit a mild upward drift in mean distance over time, consistent with the gradual increase in portfolio complexity documented in the main analysis. The persistent and substantial gap between mean and median under MinVar reflects the distributional asymmetry of that strategy's distance: a large mass of accounts with near-zero distance (pooled $\leq 1\%$ share: 63.35%) coexists with a thin but heavy right tail, pulling the mean well above the median. Under MaxRet and MaxSR the mean–median gap is narrower, indicating a less skewed distribution, consistent with their higher pooled mean distances of 28.30% and 45.09% against medians of 14.52% and 48.79%, respectively. The broad standard-deviation bands across all strategies confirm the high cross-sectional heterogeneity that motivates the portfolio-size stratification presented in the main text (Figure 8): once accounts are split by portfolio size, the MinVar mean ranges from 48.5% ($N \geq 5$) to 7.5% ($N < 5$), a spread of more than 40

percentage points that the aggregate view entirely obscures.

Figure 21 Aggregate distance over time: monthly mean and median with ±1 standard-deviation bands for each strategy, pooled across all accounts ($N = 239\,634\,756$, 72 monthly snapshots).

D Supplementary Material: Realised Portfolio Performance

This appendix supplements the realised performance analysis in Section 6. It covers the frequency of outperforming the market and the cross-sectional magnitude of beat-market margins (Section D.1), the analysis of portfolio β and market exposure (Section D.2), the association between frontier distance and risk-adjusted return (Section D.3), and the Random Forest and SHAP analysis of return and alpha predictability (Section D.4).

D.1 Frequency of Outperforming the Market

Table 2 in Section 6.1 reports realised return and alpha statistics aggregated over all account observations. As a complementary view, we report the within-snapshot share of accounts that beat the wETH and WBTC market index and summarise these per-snapshot percentages across the 72 monthly snapshots. For each strategy and each monthly snapshot we compute (i) the share of accounts whose realised 20-day return exceeds the market and (ii) the share whose CAPM alpha is strictly positive. Both numerator and denominator are restricted to accounts where the relevant return and beta values are present, so each strategy is evaluated on its own valid subpopulation. The baseline filter from Section 6.1 is applied throughout.

Table 12 reports the resulting distributions. Three patterns stand out. First, cross-strategy variation is narrow: the per-snapshot median ranges from 36.05% to 42.98% by raw return and from 36.07% to 40.23% by alpha, a spread of 6.7 and 4.2 percentage points respectively. Second, cross-snapshot variation is an order of magnitude larger: the baseline

rate alone spans 6.69% to 95.92% across snapshots, and pairwise Pearson correlations of the per-snapshot series across strategies all exceed 0.918, indicating that strategies move almost in lockstep over the calendar. Third, in 69.4% of snapshots fewer than half of baseline accounts beat the market index, reflecting the right-skewed cross section of ERC-20 returns where a long tail of underperforming tokens drags most accounts below the wETH–WBTC blend in the typical month. Together these patterns reinforce the conclusion of Section 6.1: the choice of allocation strategy has at most a marginal effect on whether an account beats the market in any given month, while the prevailing market regime explains the bulk of the variation.

The frequency view summarised in Table 12 captures only whether each account beat the market in a given month, not by how much. As a complementary magnitude view, for each monthly snapshot we compute the within-snapshot percentiles of the beat-market margin across accounts and report the median of each percentile across the 72 monthly snapshots. Two distributional metrics are reported in parallel: the raw excess return $r_s - r_{(m)}$ and the cross-sectional CAPM alpha $r_s - \beta_s\, r_{(m)}$. The per-strategy distributions are visualised in Figures 22 and 23, and selected percentiles are reported in Table 13.

Table 12 Per-snapshot frequency of beating the wETH–WBTC market index, summarised across the 72 monthly snapshots. Each cell is the unweighted median or mean across snapshots of the within-snapshot share of accounts whose realised return exceeds the market (left pair) or whose CAPM α is strictly positive (right pair).

	Beats market by return (%)		Positive CAPM α (%)	
Strategy	Median	Mean	Median	Mean
Baseline	42.98	41.55	40.23	41.69
MaxRet	36.94	40.75	36.07	40.10
MinVar	40.80	42.21	38.23	41.65
MaxSR	38.61	41.36	36.69	40.61
Equal-weight	36.05	41.79	37.20	41.24
MCap-weight	42.78	43.91	39.58	42.74

Three patterns are immediately visible. First, every strategy produces a strongly right-skewed cross-sectional distribution: the upper tail extends much further from zero than the lower tail. For the baseline portfolio the top one percent of accounts beat the market by 67.31% in a typical month while the bottom one percent return 39.87% below it, an asymmetry of 27.44 percentage points ($Q_{99} + Q_{01}$, with positive values indicating right skew). The asymmetry reproduces qualitatively across all six strategies and reflects the well-documented right skew of token-level returns on Ethereum, which dominates the cross section even after aggregation to the account level. Second, the centre of mass lies below zero for every strategy: median excess return ranges from -2.32% for MCap-weight to -4.09% for MaxRet, restating the conclusion of Table 12 as a location parameter rather than a binary frequency. Third, the cross-strategy contrasts that the figure makes most visible all live in the tails rather than at the median. MaxSR has the deepest lower tail (Q_{01} =-42.39%), reflecting the classical Markowitz problem of concentrated bets on noisy historical covariances. Equal-weight produces the narrowest distribution overall (Q_{01} =-31.41%, Q_{99} = 65.65%), trading upper-tail upside for lower-tail robustness. MCap-weight is the only strategy whose entire distribution shifts noticeably toward zero, consistent with its mechanical alignment with the market index.

The CAPM α distribution shown in Figure 23 mirrors the excess-return shape closely,

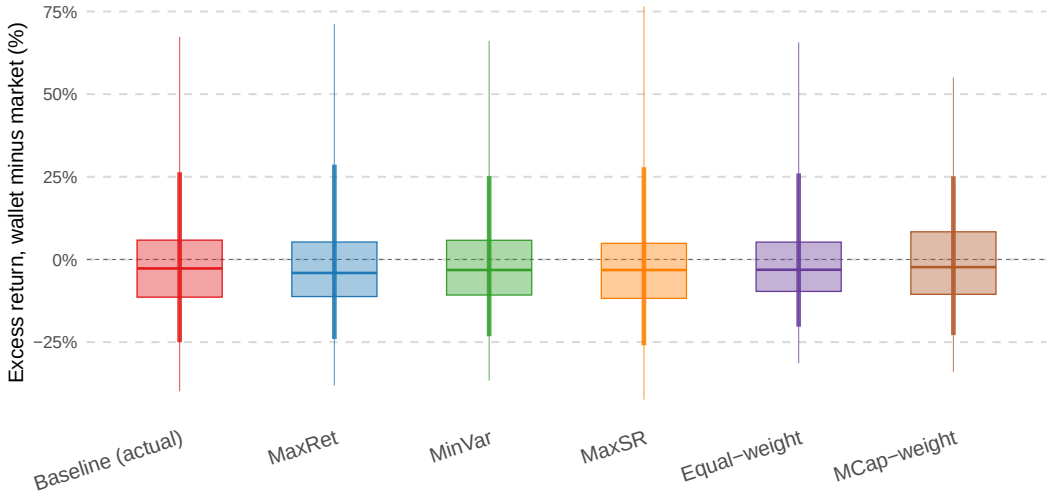

Figure 22 Cross-sectional distribution of the per-account excess return $r_s - r_{(m)}$, by strategy. Each box summarises the within-snapshot percentiles across accounts (filled box Q_{25} to Q_{75}, inner line Q_{50}, thick whisker Q_{05} to Q_{95}, thin whisker Q_{01} to Q_{99}), aggregated as the median of each percentile across the 72 monthly snapshots. The dashed line marks the market.

with slightly more pronounced tails after the β adjustment (baseline Q_{01} =-40.84% versus -39.87% on raw excess return, and $Q_{99} = 72.35\%$ versus 67.31%), indicating that accounts in both tails carry larger absolute β exposure to market moves. The cross-strategy ranking is unchanged.

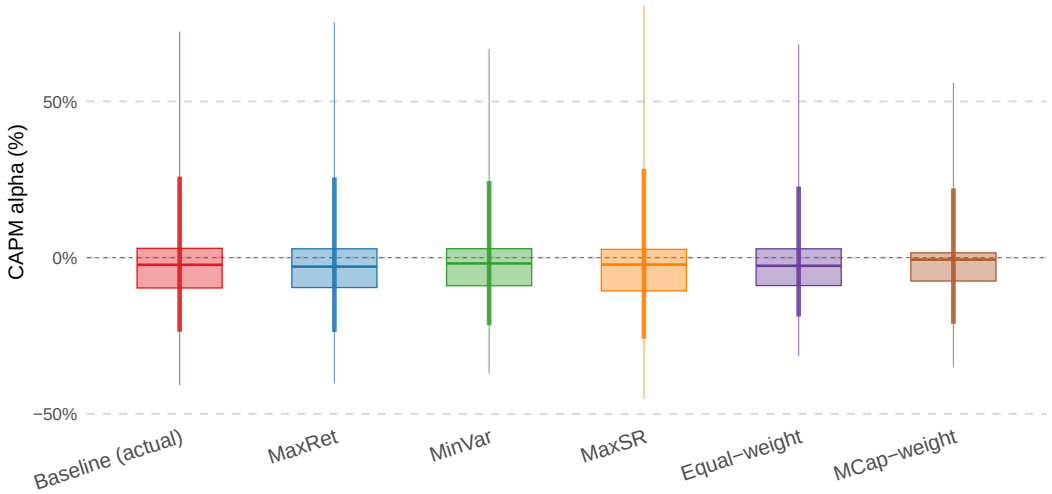

Figure 23 Distribution of the per-account CAPM α $r_s - \beta_s\, r_{(m)}$ by strategy, with quantiles aggregated as the median across all 72 monthly snapshots. Boxes are constructed as in Figure 22, with the dashed line marking zero alpha.

As a complementary temporal view, Figure 24 contrasts the between-strategy return differences with the magnitude of market-wide movements across snapshots, underscoring that allocation choice is a second-order effect relative to the market cycle.

Table 13 Distribution of the per-account beat-market margin, summarised across the 72 monthly snapshots. For each monthly snapshot we compute the listed within-snapshot percentiles across accounts and report the median of each percentile across snapshots. The *Skew* column is $Q_{99} + Q_{01}$. Positive values indicate right skew, and values are reported in percentage points.

	Excess return $r_s - r_{(m)}$ (%)				CAPM α $r_s - \beta_s\, r_{(m)}$ (%)			
Strategy	Q_{01}	Q_{50}	Q_{99}	Skew	Q_{01}	Q_{50}	Q_{99}	Skew
Baseline	−39.87	−2.71	67.31	27.44	−40.84	−2.28	72.35	31.50
MaxRet	−38.12	−4.09	71.28	33.15	−40.13	−2.86	75.31	35.17
MinVar	−36.64	−3.18	66.15	29.52	−37.23	−1.84	66.89	29.66
MaxSR	−42.39	−3.19	76.60	34.20	−45.14	−2.20	80.76	35.62
Equal-weight	−31.41	−3.12	65.65	34.25	−31.46	−2.60	68.25	36.80
MCap-weight	−34.00	−2.32	55.04	21.04	−35.06	−0.63	55.99	20.94

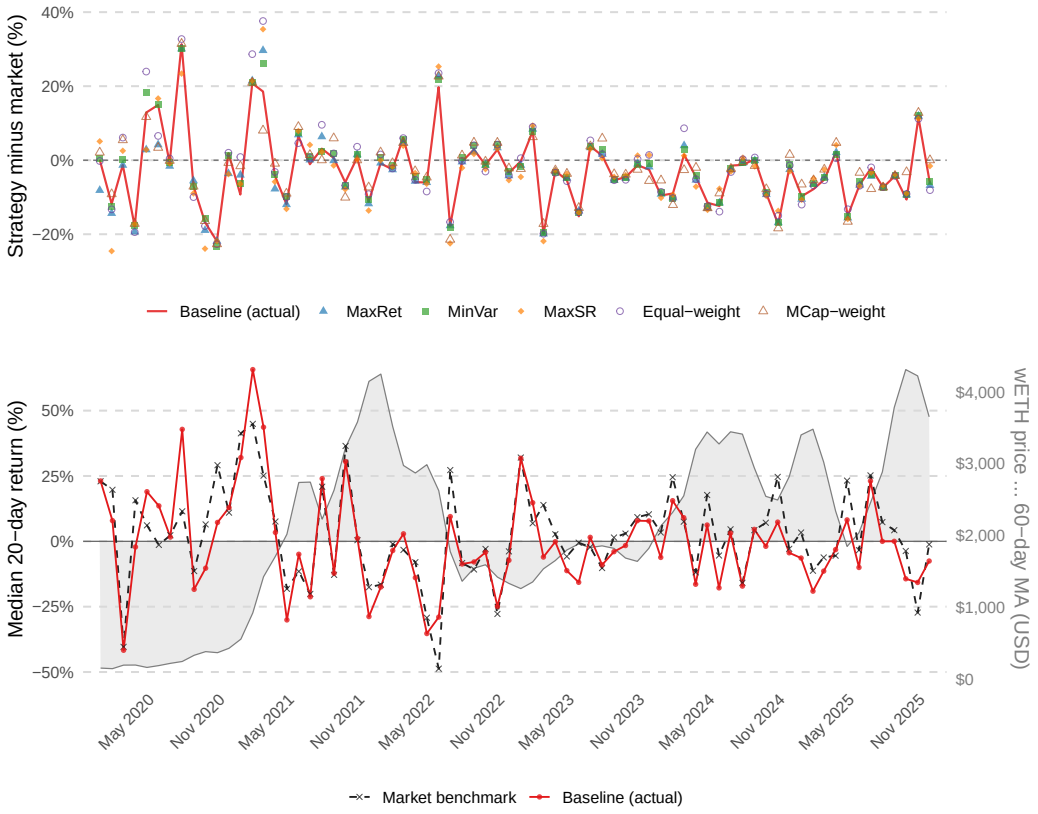

Figure 24 Returns relative to the market. Top: median $T^{(+)}$-day return minus the market benchmark per snapshot, showing that allocation choice has little effect relative to market-wide movements. Bottom: median baseline return and market benchmark overlaid on the wETH 60-day moving average price. The between-strategy spread rarely exceeds 10 pp, while the market swing spans over 100 pp.

D.2 Portfolio β and Market Exposure

Portfolio β, which captures sensitivity to the market, has a median of 0.989 across accounts (IQR 0.58–1.28) for the baseline allocation. Every optimisation strategy reduces β on average, market-cap-weighting most strongly ($\Delta\beta = -0.102$). By construction, α is orthogonal to β in CAPM (empirical $\alpha \sim \beta$: $\rho = -0.143$), so reducing market exposure provides no mechanism for improving risk-adjusted returns, and α remains negative for every strategy.

D.3 Frontier Distance and Risk-Adjusted Return

In Section 5 we showed that portfolio size is the dominant determinant of ℓ_1-distance from the efficient frontier. If MPT optimality matters in practice, we would expect a negative association between distance (d) and risk-adjusted returns (α). We find neither: Spearman rank correlations between ℓ_1-distance and baseline α are small and, if anything, positive ($\rho = +0.026$ to $+0.101$), meaning farther accounts earn slightly higher risk-adjusted returns. The raw return is higher for zero-distance accounts (-0.05%) than for the 80–100 % bucket (-6.09%), but this gap is structural rather than a payoff to optimisation: as Sections 4.3 and 5 showed, near-optimal accounts are almost exclusively two-token portfolios that mechanically land close to the frontier.

D.4 Random Forest and SHAP Details

To quantify the relative importance of account characteristics for predicting performance, we train Random Forest regressors (300 trees, depth 18, 10 M sample) on two target variables—realised $T^{(+)}$-day return (%) and cross-sectional α—for each of the six strategies. We consider two feature sets: a *baseline* model with three features (entry month, log portfolio value V, portfolio size N) and an *extended* model of six features that adds portfolio β, a binary ETH-wrapper indicator, and a binary BTC-wrapper indicator.

Feature importance.

Table 14 reports the permutation-based feature importances and R^2 for both models on the baseline strategy. The pattern is consistent across all six strategies (full results in the accompanying data release): month dominates, β is a meaningful but secondary predictor, and asset composition is negligible.

Table 14 Random Forest feature importances: permutation-based importances and out-of-sample R^2 for predicting baseline returns and α. The 6-feature model adds portfolio β, ETH-wrapper, and BTC-wrapper indicators.

	Return (%)		α	
Feature	3-feat	6-feat	3-feat	6-feat
Entry month	73%	70%	55%	55%
Portfolio β	—	22%	—	26%
Log portfolio value	26%	8%	16%	17%
Portfolio size N	1%	<1%	1%	<1%
Holds ETH wrapper	—	<1%	—	1%
Holds BTC wrapper	—	<1%	—	<1%
R^2 (test)	0.55	0.65	0.36	0.48

SHAP analysis.

We compute SHAP [32] values via TreeExplainer [31] on a 30 K subsample to assess how each feature shifts individual predictions. Figure 25 shows the SHAP summary for baseline returns. Snapshot month dominates, shifting predicted returns by ±11.4 pp on average. The portfolio β contributes ±2.7 pp, and all remaining features contribute less than ±1 pp.

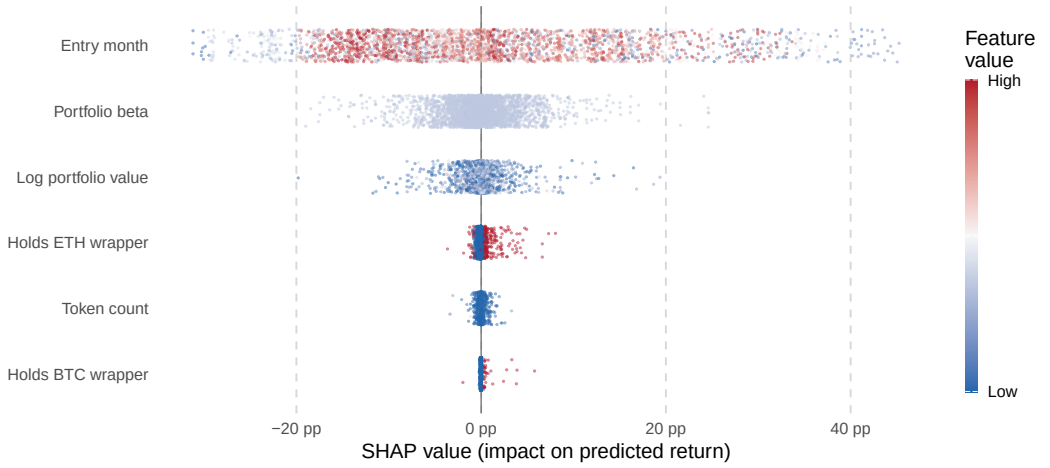

Figure 25 SHAP summary for baseline returns. Each dot represents one account-snapshot, the x-axis shows feature contribution to predicted return (pp), and colour encodes feature value (blue = low, red = high). Snapshot month dominates (±11.4 pp on average).

The SHAP dependence plot for snapshot month (Figure 26) shows predicted returns swinging by tens of percentage points across the sample period. The per-month median (red line) tracks the broader crypto market cycle. The model is not learning which accounts allocate well, rather, which months were good or bad to be in the market.

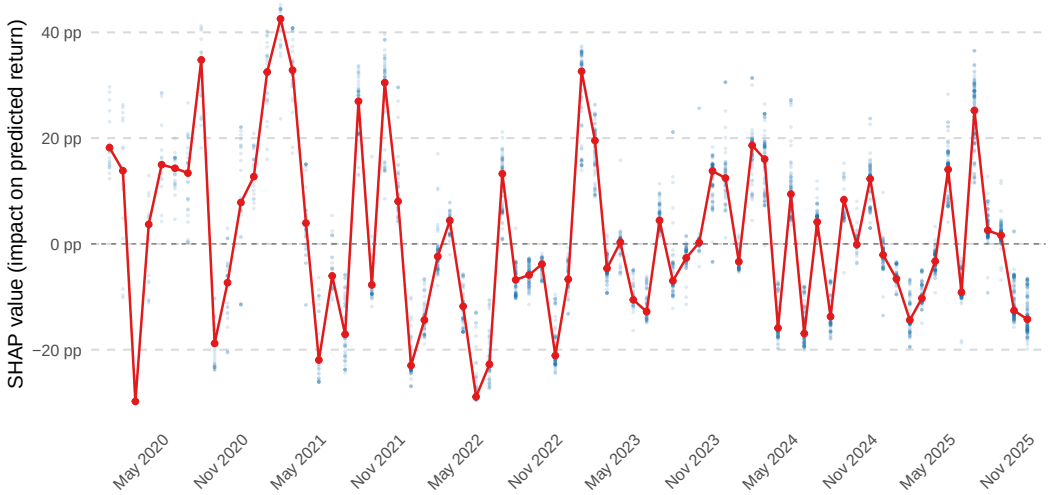

Figure 26 SHAP dependence for snapshot month. Each dot is one account-snapshot, the x-axis is the calendar month of the snapshot, and the y-axis the SHAP value (pp). The red line marks the per-month median.

3-feature vs. 6-feature comparison.

Adding β and asset-composition indicators improves R^2 by +0.07–+0.10 for returns and +0.08–+0.12 for α across all strategies (Table 15). The improvement is driven almost entirely by β: the two binary asset indicators contribute less than 1% of importance in every model. This confirms that β carries genuine predictive content beyond timing and wealth, but timing remains the dominant factor in both specifications.

Table 15 Model comparison across strategies: out-of-sample R^2 for 3-feature versus 6-feature Random Forest models.

	Return (%)			α		
Strategy	3-feat	6-feat	Δ	3-feat	6-feat	Δ
Baseline	0.55	0.65	+0.10	0.36	0.48	+0.12
MaxRet	0.51	0.60	+0.09	0.32	0.42	+0.10
MinVar	0.56	0.65	+0.09	0.37	0.47	+0.10
MaxSR	0.49	0.56	+0.07	0.29	0.38	+0.09
Equal-weight	0.58	0.65	+0.07	0.39	0.47	+0.08
MCap-weight	0.52	0.62	+0.10	0.35	0.46	+0.11

E Computational Details

This appendix describes the computational infrastructure and software stack used to produce the results presented in this paper, supplementing the methodological overview in Section 3.3. It covers the execution environment (Section E.1), the three-stage computational pipeline (Section E.2), and the software stack (Section E.3).

E.1 Execution Environment

All large-scale stages (portfolio reconstruction from raw ERC-20 logs and the downstream MPT optimisation/backtesting) were executed on the Austrian Scientific Computing (ASC) infrastructure[5], using the Vienna Scientific Cluster VSC-5[6].

All experiments were run on the default VSC-5 CPU partition `zen3_0512`[7], which provides AMD Zen 3 nodes equipped with two AMD EPYC 7713 processors (128 physical cores per node), typically with 512 GB RAM and a local 2 TB NVMe disk.

Storage model. We store all large parquet datasets on the ASC project filesystem `$DATA` (e.g., `/gpfs/data/`), which is implemented as a tiered IBM Spectrum Scale (GPFS) system with 500 TB of flash and approximately 5 PB of HDD capacity[8]. Frequently accessed files are automatically promoted to the flash tier, enabling efficient large-scale scans while keeping long-term storage feasible.

[5] https://asc.ac.at/
[6] https://asc.ac.at/systems/vsc-5/
[7] https://docs.asc.ac.at/systems/vsc5.html and https://docs.asc.ac.at/running_jobs/vsc5_queues.html
[8] https://docs.asc.ac.at/storage/where_store_data.html

E.2 Computational Pipeline

The pipeline consists of three stages: building a data lake from raw Ethereum blockchain data, reconstructing per-account portfolio snapshots, and running MPT optimisation with forward-looking backtesting.

E.2.1 Data lake infrastructure.

As a foundation to the reconstruction procedure, we utilise a data lake solution to store the raw Ethereum data. The *logs* and *traces* are the primary inputs to our data lake, organised in partitions of 10 000 Ethereum blocks in an efficiently accessible *Parquet* format. Parquet files use efficient compression and allow fast access through a scannable column-oriented storage format. This makes the format well-suited for parallel processing and provides a systematic infrastructure for processing the data.

E.2.2 Portfolio reconstruction.

The most storage-intensive phase involves extracting and parsing the complete Ethereum log history from block #0 to block #23,943,069 (Jul-30-2015 to Dec-05-2025, approximately ten years of data). We scan the complete Ethereum log history and extract transfer events for tokens matching our allow-list, yielding 1,977,546,094 transfer events across 4,847 discovered tokens, stored as per-token ledger Parquet files totalling 42.4 GB. The transfer volume is heavily concentrated: the single largest token (wETH) accounts for 698 M transfers (35.3% of the total), while the median token has only 25,030 transfers. After validation filtering (verifying ERC-20 compliance, price-data availability, sufficient trading volume, and plausible fully-diluted valuation) 4,562 tokens pass all criteria and enter the reconstruction stage.

The MapReduce approach (described in Section 3.3) proceeds in two MapReduce passes. The first pass scans the raw log partitions, decodes Transfer, Deposit, and Withdrawal events, and writes per-token shard files. If a batch exceeds its timeout, the scheduler automatically splits it in half and retries both halves, avoiding manual intervention on slow partitions. The second pass merges shards into a single sorted ledger per token, storing each transfer as two rows (a negative entry for the sender and a positive entry for the recipient) so that account balances at any block height can be recovered by cumulative summation.

Portfolio reconstruction itself reads these ledgers, aggregates balances at each of the 72 monthly snapshots via DuckDB, enriches holdings with decimals and USD prices, and writes the final snapshot Parquet. All balances for a given token–account pair are computed in a single SQL pass using conditional sums (`SUM(CASE WHEN block <= b THEN val END)`) per snapshot, rather than issuing 72 separate queries. We configure the pipeline with 48 parallel workers and a batch size of 20,000 accounts. Due to the dataset scale, reconstruction is split into two runs (2020–2022 and 2023–2025). The first run produces 477.3 M rows, the second 982.2 M rows, for a combined total of 1,459.5 M account–snapshot observations (196.6 GB compressed Parquet).

Proxy contract resolution.

A token contract may be deployed as a proxy that delegates execution to an underlying implementation. We resolve such proxies before applying the ERC-20 selector checks of Section 3.2, covering EIP-1167 minimal proxies, EIP-1967 Transparent/UUPS proxies, beacon proxies, ZeppelinOS "unstructured storage" proxies, and proxies that expose an `implementation()` method.

E.2.3 MPT optimisation and backtesting.

The optimisation stage takes the reconstructed holdings snapshots and, for each of the 72 monthly snapshots and every account with at least two tokens, solves three constrained mean–variance problems: maximise return at the account's current volatility (MaxRet), minimise variance at the account's current expected return (MinVar), and maximise the Sharpe ratio (MaxSR). Two naive benchmarks (equal-weight and market-cap-weight) are computed alongside, yielding 6 strategy variants per account. For each strategy, the pipeline evaluates realised 20-day forward returns and computes CAPM alpha and beta against a market benchmark.

The input dataset contains 239.6 M account–snapshot observations. Per-snapshot account counts range from 879,029 to 7,111,049 (median: 3,348,198), reflecting the growth of the Ethereum account base over the 2020–2025 observation window. Due to dataset scale, optimisation jobs are split by snapshot ranges and executed on 128-core nodes with 64 parallel workers per job. Price histories for all tokens in each snapshot's universe are pre-loaded into shared memory and distributed to workers via process-level initialisation, avoiding redundant I/O across the millions of per-wallet optimisations. The final output comprises 239,634,756 rows across 72 partitioned Parquet files totalling 144 GB.

E.3 Software Stack

This subsection describes the key libraries and tools underlying the pipeline described in Section E.2.

We implement the pipeline in Python 3.11 using PyArrow as the foundational columnar backend and Polars for efficient batch transformations. Parquet datasets are accessed via streaming interfaces to avoid materialising full tables in memory. For large-scale aggregation and lightweight analytical queries, we additionally employ DuckDB as an embedded analytical database (with explicitly bounded memory settings for stability). This stack provides a practical balance between scalability (HPC-ready I/O throughput), reproducibility (deterministic parquet-based artifacts), and modularity (separable reconstruction and optimisation stages).

For transparency and reproducibility, we provide the full pipeline implementation at `https://github.com/ivan-vynyavskyy/modern-portfolio-theory-in-the-crypto-wilderness`.