

# Do Whitepaper Claims Predict Market Behavior?

*Evidence from Cryptocurrency Factor Analysis*

Murad Farzulla<sup>1,2,\*</sup>

<sup>1</sup>Dissensus AI, London, UK    <sup>2</sup>King's College London, London, UK

\*Correspondence: [murad@dissensus.ai](mailto:murad@dissensus.ai)    ORCID: [0009-0002-7164-8704](https://orcid.org/0009-0002-7164-8704)

February 2026

## Abstract

This study investigates whether cryptocurrency whitepaper narratives align with empirically observed market factor structure. We construct a pipeline combining zero-shot NLP classification of 38 whitepapers across 10 semantic categories with CP tensor decomposition of hourly market data (49 assets, 17,543 timestamps). Using Procrustes rotation and Tucker's congruence coefficient ( $\phi$ ), we find weak alignment between claims and market statistics ( $\phi = 0.246, p = 0.339$ ) and between claims and latent factors ( $\phi = 0.058, p = 0.751$ ). A methodological validation comparison—statistics versus factors, both derived from market data—achieves significance ( $p < 0.001$ ), confirming the pipeline detects real structure. The null result indicates whitepaper narratives do not meaningfully predict market factor structure, with implications for narrative economics and investor decision-making. Entity-level analysis reveals specialized tokens (XMR, CRV, YFI) show stronger narrative–market correspondence than broad infrastructure tokens.

**Keywords:** Cryptocurrency, Tensor Decomposition, NLP, Factor Analysis, Procrustes Rotation, Tucker's Congruence Coefficient, Zero-Shot Classification

**JEL Codes:** G14, G12, C38, C45

**Acknowledgements.** The author acknowledges Claude (Anthropic) for assistance with pipeline development, mathematical exposition, and technical writing. All errors, omissions, and interpretive limitations remain the author's responsibility.

**Data & Code Availability.** Reproducible code and data are available at <https://github.com/studiofarzulla/tensor-defi>.

## 1 Introduction

Cryptocurrency markets present a unique laboratory for studying the relationship between narrative and price. Unlike traditional equities, where value propositions emerge gradually through earnings reports and analyst coverage, cryptocurrency projects typically articulate comprehensive visions at inception through whitepapers. These foundational documents make explicit claims about functionality, use cases, and technical architecture—claims that should, in principle, relate to how assets behave in markets.

The efficient market hypothesis suggests that asset prices reflect available information (Fama, 1970). If whitepapers constitute meaningful information about project characteristics, we might expect narrative claims to align with market behavior. Conversely, Shiller (2017) argues that “narrative economics” drives market dynamics through stories that spread virally, potentially decoupling prices from fundamentals. Recent work has begun examining this in cryptocurrency markets, with Aste (2019) documenting significant correlations between prices and sentiment across nearly two thousand cryptocurrencies. Shiller’s framework emphasizes how narratives propagate and influence aggregate behavior; we adapt this insight to test a related but distinct hypothesis about structural alignment between project-level narratives and market factor structure.

This tension motivates our central research question: *Do whitepaper claims predict market behavior?* Specifically, do the functional narratives articulated in project whitepapers align with empirically observed market factor structure? We emphasize that our methodology measures contemporaneous structural alignment between representational spaces rather than predictive relationships; the term “predict” in our research question should be understood as “exhibit structural correspondence with” rather than temporal forecasting.

We address this question through a novel methodological pipeline combining natural language processing with tensor decomposition. Our approach constructs three distinct representational spaces:

1. A **claims matrix**  $\mathbf{C} \in \mathbb{R}^{N \times K}$  derived from zero-shot classification of whitepaper text across  $K = 10$  semantic categories
2. A **market statistics matrix**  $\mathbf{S} \in \mathbb{R}^{M \times J}$  capturing  $J = 7$  financial metrics across  $M = 49$  assets
3. **Latent factors**  $\mathbf{F} \in \mathbb{R}^{M \times R}$  extracted from a high-dimensional market tensor via CP decomposition

Using Procrustes rotation and Tucker’s congruence coefficient, we then test whether these spaces align—whether assets that make similar claims exhibit similar market behavior.

Our findings reveal weak alignment across comparisons ( $\phi \leq 0.25$ ), with one notable exception: the statistics–factors comparison achieves statistical significance ( $p < 0.001$ ) despite weak magnitude, indicating systematic correspondence between market metrics and tensor-derived factors that narrative claims fail to capture. Specialized tokens (XMR, CRV, YFI, SOL) show positive alignment contributions, while DeFi infrastructure tokens (RPL, HBAR, AAVE, SUSHI) exhibit the largest narrative–market divergence. This result is robust to temporal variation, subsample perturbation, decomposition method (CP vs Tucker), and rank selection.

**Contributions.** This paper makes four contributions: (1) we introduce a reproducible pipeline for comparing textual and market representational spaces in cryptocurrency research, (2) we provide detailed methodological exposition of tensor decomposition for factor extraction, (3) we deliver rigorous empirical evidence on the (mis)alignment between project narratives and market behavior, and (4) we demonstrate that null results in this domain constitute valid findings with implications for narrative economics.

The remainder of this paper is organized as follows. Section 2 reviews related work. Section 3 describes our data sources. Section 4 provides detailed methodology with full mathematical exposition. Section 5 presents comprehensive results including sensitivity analyses. Section 6 interprets findings, and Section 7 concludes.

## 2 Related Work

### 2.1 Cryptocurrency Narratives and Sentiment

Research on cryptocurrency narratives spans social media analysis, whitepaper studies, and sentiment measurement. [Chen et al. \(2019\)](#) demonstrate that machine learning models trained on engineered sample dimensions can predict Bitcoin price movements, while [Ante \(2023\)](#) find that Elon Musk’s tweets generate significant abnormal returns for mentioned cryptocurrencies. [Haykir and Yağlı \(2022\)](#) document speculative bubbles driven by narrative contagion across crypto assets. [Liu and Tsyvinski \(2021\)](#) establish that cryptocurrency returns exhibit momentum, size effects, and exposure to market-wide factors distinct from traditional assets.

Whitepaper analysis has received growing attention as a signal of project quality. [Howell et al. \(2020\)](#) examine ICO whitepaper quality as a signal of project legitimacy, finding that technical depth correlates with fundraising success. [Fisch \(2019\)](#) show that whitepaper informativeness predicts ICO outcomes. [Adhami et al. \(2018\)](#) provide the first comprehensive empirical description of ICO determinants, finding that code availability, token presale, and utility design predict success. More recent work has applied sophisticated NLP to whitepaper content: [Thewissen et al. \(2022\)](#) use topic modeling on 5,210 whitepapers to show that technical feature topics predict success while whitepaper informativeness diminishes post-listing. Critically, [Momtaz \(2021\)](#) demonstrates that issuers systematically exaggerate whitepaper claims—exaggeration raises funds short-term but causes token depreciation and platform failure long-term. [Samieifar and Baur \(2021\)](#) find that whitepaper length and complexity correlate positively with funds raised, supporting signaling interpretations. [Florysiak and Schandlbauer \(2022\)](#) show that while expert ratings initially “jam” whitepaper signals, post-listing returns are better predicted by whitepaper content than analyst assessments.

However, these studies focus on cross-sectional prediction at issuance rather than ongoing alignment between narrative and market behavior. Indeed, [Suriano et al. \(2025\)](#) find that clustering cryptocurrencies by whitepaper content does not yield significant differences in time series dynamics, suggesting narratives may matter for initial fundraising but not long-term price behavior. Recent high-profile failures illustrate the gap between whitepaper claims and realized outcomes: [Briola et al. \(2022\)](#) analyze the Terra-Luna collapse, demonstrating how algorithmic stablecoin mechanisms failed despite elaborate whitepaper specifications, while [Vidal-Tomás et al. \(2023\)](#) document how centralized exchange fragility propagates through cryptocurrency markets. Our work extends this by testing whether claims align with ongoing market factor structure.

**Narrative Economics Framework.** [Shiller \(2017\)](#) provides the theoretical foundation for studying how narratives shape economic outcomes. This builds on foundational work in behavioral finance: [Baker and Wurgler \(2006\)](#) construct the canonical investor sentiment index, demonstrating that sentiment-driven mispricing is strongest for speculative, hard-to-value stocks—a category that includes most cryptocurrencies. [Baker and Wurgler \(2007\)](#) survey evidence on how sentiment affects both aggregate returns and cross-sectional pricing. [Tetlock \(2007\)](#) shows that media pessimism predicts market downturns and high trading volume, establishing the textual analysis framework for investor sentiment. Applied to cryptocurrency, this framework suggests that project narratives—embodied in whitepapers—should in-

fluence investor beliefs and thus market prices. Our study tests this prediction empirically, finding weak support for narrative-market coupling. This complements information-theoretic approaches (Keskin and Aste, 2020) that find nonlinear causality between social sentiment and cryptocurrency returns.

**Natural Language Processing in Finance.** The application of NLP to financial text has grown substantially, with transformer-based models enabling sophisticated document analysis (Loughran and McDonald, 2020). Loughran and McDonald (2011) demonstrate that general sentiment dictionaries misclassify financial text, motivating domain-specific approaches. Kearney and Liu (2014) provide a comprehensive methodological survey of textual sentiment methods and models. Domain-specific pre-training has proven effective: Araci (2019) introduce FinBERT, outperforming Loughran-McDonald dictionaries on financial sentiment, while Huang et al. (2023) show similar gains on analyst reports. Mischev et al. (2020) benchmark lexicons against transformers (BERT, FinBERT) across financial datasets, finding transformer superiority for nuanced sentiment. Zero-shot classification, as employed here, allows categorization without domain-specific training data (Lewis et al., 2020). Bartolucci et al. (2020) demonstrate that sentiment extracted from developer communications (GitHub) predicts cryptocurrency price movements, validating NLP approaches in this domain. Our use of BART-MNLI represents a middle ground—leveraging powerful pre-trained representations while acknowledging potential limitations in cryptocurrency-specific semantics.

**Cryptocurrency Market Microstructure.** Beyond sentiment, cryptocurrency markets exhibit distinctive microstructure features: 24/7 trading, global fragmentation across exchanges, varying levels of market manipulation, and high correlation with Bitcoin. These features may dominate narrative effects. Pappalardo et al. (2018) document inefficiencies in Bitcoin’s peer-to-peer network that diverge from whitepaper ideals of decentralization. Farzulla (2025) demonstrate that infrastructure disruption events generate significantly larger volatility responses than regulatory announcements, suggesting market participants weight technical fundamentals over policy uncertainty. This asymmetry—where operational failures matter more than regulatory shifts—implies that whitepaper claims about technical capabilities may warrant particular attention in volatility modeling. Our inclusion of multiple market statistics (volatility, liquidity, drawdown) attempts to capture this microstructure, but residual factors may remain.

## 2.2 Factor Models in Cryptocurrency

Traditional asset pricing employs factor models to explain cross-sectional return variation. Fama (1970) established the theoretical foundation for efficient markets and factor-based returns. Fama and French (1993) introduced the three-factor model for equities; analogous developments in cryptocurrency have emerged more recently. Livan et al. (2011) demonstrate how random matrix theory can distinguish signal from noise in financial correlation matrices—a perspective we extend to narrative-factor comparisons. Caccioli et al. (2018) provide a comprehensive review of network-based approaches to financial systemic risk, complementing factor-based perspectives with topological analysis. More recently, Aste (2025) establishes the theoretical and generative foundations for information filtering networks, offering principled methods for extracting sparse dependency structures from high-dimensional financial data that may inform future network-aware factor models.

Liu et al. (2022) establish the foundational three-factor model for cryptocurrency returns: market, size, and momentum. Liu and Tsyvinski (2021) extend this work, finding that these factors explain substantial cross-sectional variation analogous to Fama-French factors for equities. Bianchi and Babiak (2021) apply Instrumented PCA (IPCA) to show that time-varying factor loadings outperform observable risk factors. Dobrynskaya (2020) extends crypto CAPM to include downside beta as a fourth factor, finding significant cross-sectional premia across 1,700 coins. Bhambhwani et al. (2019) introduce

blockchain-native factors—computing power, network size—as procyclical pricing factors with positive risk premia.

Our tensor decomposition implicitly captures similar factors—Factor 1 (dominated by Bitcoin) resembles the market factor, while Factor 2 may capture size or sector effects. These systematic factors may dominate any narrative-based signal, explaining why whitepaper claims fail to predict factor structure.

**Multi-Way Data in Finance.** Financial data naturally exhibits multi-way structure: assets  $\times$  time  $\times$  features. While matrix methods (PCA, factor analysis) collapse this structure, tensor decomposition preserves it. Our work contributes to the emerging literature applying tensor methods to financial data, demonstrating both their utility (interpretable factors) and limitations (modest alignment despite high explanatory power).

### 2.3 Tensor Methods in Finance

Tensor decomposition provides a natural framework for analyzing multi-way financial data. [Kolda and Bader \(2009\)](#) provide a comprehensive review of tensor decomposition methods, establishing the theoretical foundations for CP and Tucker decomposition. Recent work has applied these methods to financial time series: [Chen et al. \(2022\)](#) develop tensor factor models for high-dimensional time series, introducing TIPUP/TOPUP estimators with explicit applications to economics and finance. [Wang et al. \(2022\)](#) apply Tucker decomposition to high-dimensional vector autoregression, demonstrating improved performance over matrix-based approaches for multivariate time series. [Han et al. \(2024\)](#) develop CP factor models for dynamic tensors, providing uncorrelated latent factors directly applicable to asset pricing. [Fan et al. \(2013\)](#) introduce the POET estimator for high-dimensional covariance estimation with factor structure. [Briola et al. \(2026\)](#) extend PCA to incorporate graph structure, regularizing principal components with network topology—an approach that could inform future tensor decompositions where asset relationships carry meaningful structural information beyond covariance.

CP (CANDECOMP/PARAFAC) decomposition decomposes a tensor into rank-one components, extracting interpretable latent factors ([Harshman, 1970](#)). For market data structured as (time  $\times$  asset  $\times$  feature), CP decomposition yields asset-level factor loadings analogous to principal component analysis but preserving multi-way structure. Tucker decomposition offers an alternative with mode-specific ranks and a core tensor capturing interactions.

### 2.4 Factor Comparison Methods

Comparing factor structures across studies or datasets requires methods that account for rotational indeterminacy. Procrustes rotation ([Schönemann, 1966](#)) finds the optimal orthogonal transformation aligning one factor matrix to another. [Brokken \(1983\)](#) develops orthogonal Procrustes rotation that directly maximizes congruence, the approach we employ. Tucker’s congruence coefficient ( $\phi$ ) then measures similarity between aligned factors ([Tucker, 1951](#); [Lorenzo-Seva and ten Berge, 2006](#)).

[Korth and Tucker \(1975\)](#) establish null distributions for congruence coefficients from simulated data, essential for understanding what constitutes statistically meaningful alignment. [Paunonen \(1997\)](#) examines the distribution of factor congruence under chance conditions after Procrustes rotation, informing our significance testing. [Lorenzo-Seva and ten Berge \(2006\)](#) establish interpretation thresholds:  $|\phi| \geq 0.95$  indicates factor equivalence,  $|\phi| \geq 0.85$  indicates fair similarity,  $|\phi| \geq 0.65$  indicates moderate similarity, and  $|\phi| < 0.65$  indicates weak or no similarity. These thresholds guide our interpretation of narrative-market alignment.

### 3 Data

#### 3.1 Market Data

We collect hourly OHLCV (open, high, low, close, volume) data via the Binance exchange API (accessed through the CCXT library) for 49 cryptocurrency assets spanning January 1, 2023 to December 31, 2024, yielding 17,543 timestamps per asset. Asset selection follows liquidity and data availability criteria, including major cryptocurrencies (BTC, ETH) and a diverse set of DeFi, infrastructure, and utility tokens.

Table 1 summarizes the dataset dimensions.

Table 1: Dataset Summary

Dimension	Value
Assets (market data)	49
Assets (whitepapers)	38
Assets (common intersection)	37
Time period	Jan 2023 – Dec 2024
Timestamps (hourly)	17,543
Market features (OHLCV)	5
Derived statistics	7
Narrative categories	10

#### 3.2 Whitepaper Corpus

We collect whitepapers for 38 assets where official foundational documents are publicly available, including AAVE, ADA, ALGO, AR, ARB, ATOM, AVAX, BTC, COMP, DOT, ENS, ETH, FIL, GRT, ICP, LINK, MKR, NEAR, SC, SOL, STORJ, UNI, XMR, and ZEC (see Table 2 for the full asset list). Documents include original whitepapers, consensus papers, protocol specifications, and technical documentation. PDF text is extracted using PyPDF2 with sentence-level tokenization; for assets without extractable PDFs, we use official documentation in markdown format. Table 2 summarizes corpus statistics for selected assets.

Table 2: Whitepaper Corpus Statistics (Selected)

Asset	Pages	Year	Type
ZEC	229	2020	Protocol Spec
STORJ	90	2018	Storage WP
ADA	48	2020	Consensus
ICP	45	2021	Tech Overview
LINK	38	2017	Oracle WP
FIL	36	2017	Tech Report
ETH	36	2014	Original WP
SOL	32	2018	Original WP
NEAR	23	2020	Sharding
MKR	21	2017	Stablecoin
XMR	20	2013	CryptoNote

+ 27 additional documents (see Appendix)

Table 3 summarizes data availability across the pipeline. The intersection of whitepaper ( $n = 38$ ) and market data ( $n = 49$ ) yields 37 common assets for alignment analysis. AR is excluded due to insufficient market data coverage.

Table 3: Data Flow and Asset Coverage

Data Source	Assets	Notes
Whitepaper corpus	38	Technical documents
Market data (Binance)	49	2-year hourly OHLCV
Tensor factors (CP)	49	Rank-2 decomposition
$NLP \cap Market$	37	Primary analysis sample
Excluded	1	AR (market data gap)
Market-only	12	No whitepaper available

## 4 Methodology

Our pipeline proceeds in five stages: (1) tensor construction, (2) tensor decomposition, (3) NLP claims extraction, (4) market statistics computation, and (5) Procrustes alignment with congruence testing.

### 4.1 Tensor Construction

**Definition 4.1** (Market Tensor). A market tensor  $\mathcal{X} \in \mathbb{R}^{T \times V \times A \times F}$  is a 4-way array with modes:

- Time ( $T = 17,543$  hourly timestamps)
- Venue ( $V = 1$ , Binance)
- Asset ( $A = 49$  cryptocurrencies)
- Feature ( $F = 5$ , OHLCV)

With a single venue, the effective structure is 3-way:  $\mathcal{X} \in \mathbb{R}^{T \times A \times F}$ . Each entry  $x_{taf}$  represents the value of feature  $f$  for asset  $a$  at time  $t$ . Prior to decomposition, we z-normalize each feature slice across

both assets and time (i.e., each  $\mathcal{X}_{::f}$  has zero mean and unit variance), ensuring that scale differences across OHLCV features do not dominate the factor structure.

## 4.2 Tensor Decomposition

### 4.2.1 CP Decomposition

**Definition 4.2** (CP Decomposition). The CANDECOMP/PARAFAC (CP) decomposition approximates a tensor as a sum of rank-one tensors:

$$\mathcal{X} \approx \sum_{r=1}^R \lambda_r \mathbf{a}_r \circ \mathbf{b}_r \circ \mathbf{c}_r \quad (1)$$

where  $\circ$  denotes outer product,  $\lambda_r$  are weights, and  $\mathbf{a}_r \in \mathbb{R}^T$ ,  $\mathbf{b}_r \in \mathbb{R}^A$ ,  $\mathbf{c}_r \in \mathbb{R}^F$  are mode-specific factor vectors.

The factor matrices are:

$$\mathbf{A} = [\mathbf{a}_1 | \cdots | \mathbf{a}_R] \in \mathbb{R}^{T \times R} \quad (\text{time factors}) \quad (2)$$

$$\mathbf{B} = [\mathbf{b}_1 | \cdots | \mathbf{b}_R] \in \mathbb{R}^{A \times R} \quad (\text{asset factors}) \quad (3)$$

$$\mathbf{C} = [\mathbf{c}_1 | \cdots | \mathbf{c}_R] \in \mathbb{R}^{F \times R} \quad (\text{feature factors}) \quad (4)$$

The asset factor matrix  $\mathbf{B}$  provides latent loadings for alignment testing.

### 4.2.2 Alternating Least Squares

CP decomposition is computed via alternating least squares (ALS):

#### Algorithm 1: CP-ALS

- 1: Initialize  $\mathbf{A}$ ,  $\mathbf{B}$ ,  $\mathbf{C}$  randomly
- 2: **repeat**
- 3:    $\mathbf{A} \leftarrow \mathbf{X}_{(1)}(\mathbf{C} \odot \mathbf{B})(\mathbf{C}^\top \mathbf{C} * \mathbf{B}^\top \mathbf{B})^\dagger$
- 4:    $\mathbf{B} \leftarrow \mathbf{X}_{(2)}(\mathbf{C} \odot \mathbf{A})(\mathbf{C}^\top \mathbf{C} * \mathbf{A}^\top \mathbf{A})^\dagger$
- 5:    $\mathbf{C} \leftarrow \mathbf{X}_{(3)}(\mathbf{B} \odot \mathbf{A})(\mathbf{B}^\top \mathbf{B} * \mathbf{A}^\top \mathbf{A})^\dagger$
- 6: **until** convergence

where  $\mathbf{X}_{(n)}$  is mode- $n$  matricization,  $\odot$  is Khatri-Rao product,  $*$  is Hadamard product, and  $\dagger$  denotes pseudoinverse.

### 4.2.3 Rank Selection

We select rank  $R$  to achieve target explained variance:

$$\text{EV}(R) = 1 - \frac{\|\mathcal{X} - \hat{\mathcal{X}}_R\|_F^2}{\|\mathcal{X} - \bar{x}\|_F^2} \quad (5)$$

With target  $\text{EV} \geq 0.90$ , we obtain  $R = 2$  ( $\text{EV} = 92.45\%$ ).<sup>1</sup>

<sup>1</sup>Standard PCA on mode-2 (asset) matricization yields nearly identical alignment ( $\phi = 0.058$ , vs CP  $\phi = 0.058$ ), confirming that preserving tensor structure does not alter our findings. We retain the tensor framework for interpretability of multi-way temporal dynamics and consistency with recent financial tensor methods (Chen et al., 2022).

#### 4.2.4 Tucker Decomposition

For robustness, we also implement Tucker decomposition:

$$\mathcal{X} \approx \mathcal{G} \times_1 \mathbf{A} \times_2 \mathbf{B} \times_3 \mathbf{C} \quad (6)$$

where  $\mathcal{G} \in \mathbb{R}^{R_1 \times R_2 \times R_3}$  is the core tensor and  $\times_n$  denotes mode- $n$  product.

### 4.3 NLP Claims Extraction

#### 4.3.1 Zero-Shot Classification

We employ BART-large-MNLI (Lewis et al., 2020) for zero-shot text classification via the Hugging-Face Transformers library.<sup>2</sup> Whitepapers are segmented into 500-word chunks ( $n = 2,056$  across the initial 24-asset corpus); this chunk size balances computational efficiency with context preservation. The corpus was subsequently expanded to 38 assets as additional whitepapers were identified (see Table 2). Each chunk is classified against ten domain-relevant categories. Zero-shot classification follows the entailment approach of Yin et al. (2019), constructing hypotheses of the form “This text is about [category]” for each candidate label.

Given a text segment  $t$  and candidate labels  $\{l_1, \dots, l_K\}$ , the model computes:

$$P(l_k|t) = \frac{\exp(s_k)}{\sum_{j=1}^K \exp(s_j)} \quad (7)$$

where  $s_k$  is the entailment score for label  $l_k$ . Rather than argmax classification, we compute probability-weighted category profiles for each entity, providing smoother estimates that account for semantic ambiguity in technical prose.

#### 4.3.2 Semantic Taxonomy

Our taxonomy comprises  $K = 10$  categories capturing core blockchain functionality (Table 4).

Table 4: Semantic Category Taxonomy

Category	Description
Store of Value	Digital gold, inflation hedge, wealth preservation
Medium of Exchange	Payment system, transactions, currency
Smart Contracts	Programmable contracts, automation, trustless execution
Decentralized Finance	Lending, borrowing, yield, liquidity provision
Governance	Voting, DAOs, community decision-making
Scalability	High throughput, low latency, Layer 2 solutions
Privacy	Anonymous transactions, zero-knowledge proofs
Interoperability	Cross-chain communication, bridges, multi-chain
Data Storage	Decentralized storage, file systems, permanence
Oracle Services	External data feeds, real-world information

<sup>2</sup>Model: `facebook/bart-large-mnli`. BART-large fine-tuned on Multi-Genre Natural Language Inference (MNLI; Williams et al. 2018) for natural language inference.

### 4.3.3 Model Validation

We assess classification reliability through inter-model agreement using DeBERTa-v3 (He et al., 2021) as an alternative classifier.<sup>3</sup> On a random sample of 200 chunks, exact top-1 agreement is 32% (Cohen’s  $\kappa = 0.14$ ), reflecting known sensitivity of zero-shot NLI to model-specific category boundaries. However, relaxed agreement—where the alternative model’s top prediction appears in the primary model’s top-3—reaches 67%, suggesting models capture similar semantic neighborhoods with different decision thresholds. Recent advances in financial sentiment estimation using large language models (Kirtac and Germano, 2025) suggest alternative approaches for future work.

Bootstrap 95% confidence intervals on aggregate category proportions (1,000 resamples) yield tight bounds: Smart Contracts 26.3–30.1%, Scalability 18.4–22.1%, Governance 14.7–17.9%, indicating stable estimates at the corpus level despite chunk-level uncertainty.

**Multi-Method Validation.** To further assess classification robustness, we implement three independent methods with distinct inductive biases: (1) BART-MNLI (NLI-based entailment), (2) sentence embeddings using all-mpnet-base-v2 (Reimers and Gurevych, 2019) with cosine similarity to category descriptions, and (3) Minstral-3 3B, a local language model via structured JSON prompting. Table 5 reports pairwise correlations across the 38-asset expanded corpus.

Table 5: Multi-Method Classification Agreement

Method Pair	Pearson $r$	Spearman $\rho$
BART-NLI vs Embedding	0.103	0.111
BART-NLI vs LLM	0.413	0.457
Embedding vs LLM	0.392	0.384
<b>Mean pairwise</b>	<b>0.302</b>	<b>0.317</b>

The LLM-based classifier exhibits moderate correlation with both other methods ( $r \approx 0.4$ ), while BART and embedding methods show weaker agreement ( $r = 0.10$ ), suggesting distinct inductive biases. Discretized Fleiss’ Kappa ( $\kappa = 0.045$ ) indicates slight but positive inter-rater agreement above chance.

Notably, per-category agreement varies substantially (see Appendix G for full heatmap): categories with clear linguistic markers show strong convergence (medium\_of\_exchange  $\bar{r} = 0.65$ , DeFi  $\bar{r} = 0.63$ , privacy  $\bar{r} = 0.54$ ), while abstract concepts show weaker agreement (smart\_contracts  $\bar{r} = 0.24$ , store\_of\_value  $\bar{r} = 0.38$ ). This pattern suggests that some functional categories are more robustly identifiable across methods than others—a finding with implications for taxonomy design in cryptocurrency NLP research. Figure 1 visualizes these cross-method patterns.

<sup>3</sup>Model: cross-encoder/nli-deberta-v3-small.

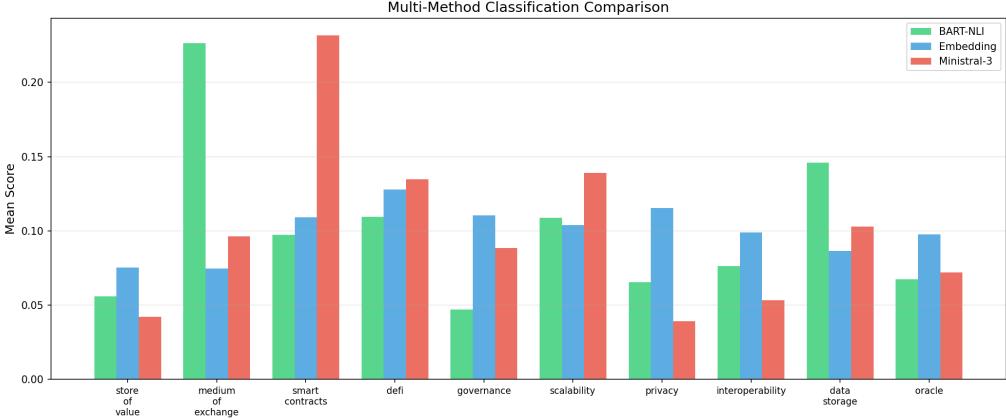


Figure 1: Multi-method classification comparison across 10 semantic categories. Bar heights represent mean classification scores for BART-NLI (primary), sentence embeddings, and Minstral-3 LLM. Categories with high cross-method agreement (medium\_of\_exchange, DeFi, privacy) show consistent relative rankings; categories with weak agreement (smart\_contracts, store\_of\_value) exhibit larger inter-method variance.

**Data Quality Filtering.** Seven assets (ALGO, AXS, BAND, DOT, RPL, SUSHI, YFI) contain fewer than 10 text chunks, yielding unreliable classification estimates. Excluding these low-data assets ( $n = 31$  reliable), the embedding–LLM correlation improves substantially ( $r = 0.52$ ), suggesting that classification disagreement partially reflects data sparsity rather than fundamental method divergence.

#### 4.3.4 Aggregation

For each asset  $n$ , we aggregate probability-weighted classification scores across text chunks:

$$c_{nk} = \frac{1}{|T_n|} \sum_{t \in T_n} P(l_k | t) \quad (8)$$

yielding claims matrix  $\mathbf{C} \in \mathbb{R}^{N \times K}$ .

#### 4.3.5 Institutional Corpus (Snapshot)

To test whether broader institutional narratives improve alignment beyond static whitepapers, we construct a supplementary corpus capturing evolving utility positioning. This corpus includes official documentation, foundation updates, governance forum posts, and project blogs. For this analysis we use a February 2025 snapshot covering 14 assets (AAVE, ADA, AR, ARB, AVAX, BTC, ETH, FIL, IMX, LINK, MKR, OP, SOL, UNI), yielding 627 documents and 367,071 words. Of these, 12 have sufficient market data coverage for rolling window alignment analysis (AR and IMX excluded due to incomplete hourly data). We classify these documents using the same taxonomy to build an institutional claims matrix. With a single snapshot, the claims matrix is held constant across rolling market windows; future work extends this to multi-period narrative drift analysis.

#### 4.4 Market Statistics

We compute seven summary statistics for each asset, then z-normalize cross-sectionally (across assets) to ensure comparability:

1. **Mean return:**  $\bar{r}_a = \frac{1}{T} \sum_t r_{at}$

2. **Volatility:**  $\sigma_a = \sqrt{\frac{1}{T-1} \sum_t (r_{at} - \bar{r}_a)^2}$
3. **Sharpe ratio**<sup>4</sup>:  $SR_a = \frac{\bar{r}_a}{\sigma_a} \cdot \sqrt{252 \cdot 24}$
4. **Max drawdown:**  $MDD_a = \min_t \frac{P_t - \max_{s \leq t} P_s}{\max_{s \leq t} P_s}$
5. **Avg volume:**  $\bar{V}_a = \frac{1}{T} \sum_t V_{at}$
6. **Vol-of-vol:**  $\sigma_{\sigma,a}$  (rolling volatility std)
7. **Trend:**  $\beta_a$  from  $P_t = \alpha + \beta t + \varepsilon$

This yields statistics matrix  $\mathbf{S} \in \mathbb{R}^{M \times 7}$ .

## 4.5 Procrustes Alignment

### 4.5.1 Problem Formulation

**Definition 4.3** (Orthogonal Procrustes Problem). Given matrices  $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times p}$ , find orthogonal  $\mathbf{Q} \in \mathbb{R}^{p \times p}$  minimizing:

$$\min_{\mathbf{Q}^\top \mathbf{Q} = \mathbf{I}} \|\mathbf{AQ} - \mathbf{B}\|_F^2 \quad (9)$$

### 4.5.2 SVD Solution

**Theorem 4.1** (Schönemann 1966). *The optimal rotation is  $\mathbf{Q}^* = \mathbf{V}\mathbf{\Sigma}\mathbf{V}^\top$  where  $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top = SVD(\mathbf{A}^\top \mathbf{B})$ .*

*Proof.* See Appendix A. ■

### 4.5.3 Dimension Handling

When comparing matrices of different column dimensions, we zero-pad the smaller matrix to match dimensions before alignment. For example, comparing the 10-dimensional claims matrix to the 2-dimensional factor matrix requires padding the factor matrix with 8 zero columns. This approach preserves all information in both matrices but introduces a potential downward bias in  $\phi$ : zero-padded dimensions contribute nothing to the numerator but may affect the Procrustes rotation. We prefer this conservative approach to the alternative of PCA-reducing the claims matrix, which would discard potentially informative narrative dimensions.

## 4.6 Tucker's Congruence Coefficient

**Definition 4.4** (Tucker's  $\phi$ ). The congruence coefficient between vectors  $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$  is:

$$\phi(\mathbf{x}, \mathbf{y}) = \frac{\sum_{i=1}^n x_i y_i}{\sqrt{\sum_{i=1}^n x_i^2 \cdot \sum_{i=1}^n y_i^2}} \quad (10)$$

This equals cosine similarity without mean-centering, appropriate for factor comparison where sign and magnitude both carry meaning.

### 4.6.1 Matrix Congruence

For matrices  $\mathbf{A}, \mathbf{B}$  after Procrustes alignment, we compute per-column  $\phi$  values and report mean absolute  $\phi$ :

$$\bar{\phi} = \frac{1}{p} \sum_{j=1}^p |\phi(\mathbf{a}_j, \mathbf{b}_j)| \quad (11)$$

---

<sup>4</sup>We use 252 trading days following traditional finance conventions for comparability with academic literature, though cryptocurrency markets operate continuously.

#### 4.6.2 Interpretation Thresholds

Following [Lorenzo-Seva and ten Berge \(2006\)](#)<sup>5</sup>:

- $|\phi| \geq 0.95$ : Factor equivalence
- $|\phi| \geq 0.85$ : Fair similarity
- $|\phi| \geq 0.65$ : Moderate similarity
- $|\phi| < 0.65$ : Weak/no similarity

### 4.7 Statistical Inference

#### 4.7.1 Power Considerations

With  $n = 37$  common entities, statistical power to detect alignment is limited. Monte Carlo simulation (500 iterations per effect size, 200 permutations each) reveals approximate power at  $\alpha = 0.05$ :

True $\phi$	Power
0.30	14%
0.50	45%
0.65	70%
0.70	90%

This analysis indicates our study is adequately powered ( $>80\%$ ) to detect only strong alignment ( $\phi \geq 0.70$ ). The “moderate similarity” threshold of  $\phi = 0.65$  falls near our detection limit. Consequently, our null findings should be interpreted cautiously: we can confidently reject strong alignment but cannot distinguish weak alignment ( $\phi \approx 0.3$ ) from no alignment. The non-significant p-values for claims-based comparisons are consistent with both no alignment and insufficient power to detect weak effects.

#### 4.7.2 Permutation Test

We assess significance via one-sided permutation test (implemented in Python using `scipy.stats` and `numpy`):

1. Compute observed  $\phi^*$
2. For  $b = 1, \dots, B$  permutations: permute rows of  $\mathbf{B}$ , compute  $\phi^{(b)}$
3.  $p\text{-value} = \frac{1}{B} \sum_{b=1}^B \mathbf{1}[\phi^{(b)} \geq \phi^*]$  (one-sided, testing  $H_0: \phi \leq \phi_{\text{random}}$ )

#### 4.7.3 Bootstrap Confidence Intervals

We construct 95% CIs via percentile bootstrap ( $B = 1000$  resamples). However, bootstrap resampling with replacement on small samples ( $n = 37$ ) exhibits known pathologies when combined with Procrustes-based alignment: duplicate entities in resampled data artificially inflate  $\phi$  by increasing effective weights on well-aligned pairs. Our bootstrap distributions show substantial upward bias (bootstrap mean exceeds point estimate by 29% for claims–statistics, 37% for claims–factors), with moderate right-skewness (skewness  $\approx 0.45$ ). Consequently, percentile CIs may be conservative for upper bounds but unreliable for lower bounds. Given this bias, reported confidence intervals should be interpreted as indicative rather than precise. We report these intervals for completeness while acknowledging their limitations.

<sup>5</sup>These thresholds were developed for comparing factor solutions derived from similar data (e.g., across samples or rotation methods). Their application to heterogeneous spaces—comparing NLP-derived claims to market-derived factors—extends beyond the original validation context. We apply them as rough benchmarks while acknowledging this limitation.

#### 4.7.4 Reproducibility

All stochastic procedures use fixed random seeds (seed = 42 for CP-ALS initialization, tensor operations, and permutation tests) to ensure reproducibility. Permutation tests use  $B = 1000$  iterations; bootstrap procedures use  $B = 1000$  resamples. Implementation uses Python 3.11+ with NumPy, SciPy, TensorFlow (CP-ALS), scikit-learn (CCA, PLS), and Hugging Face Transformers (BART-MNLI). Code and data are available at <https://github.com/studiofarzulla/tensor-defi>.

## 5 Results

### 5.1 Tensor Decomposition

CP decomposition with rank 2 explains 92.45% of tensor variance. Figure 2 visualizes a cross-sectional slice of the market tensor.

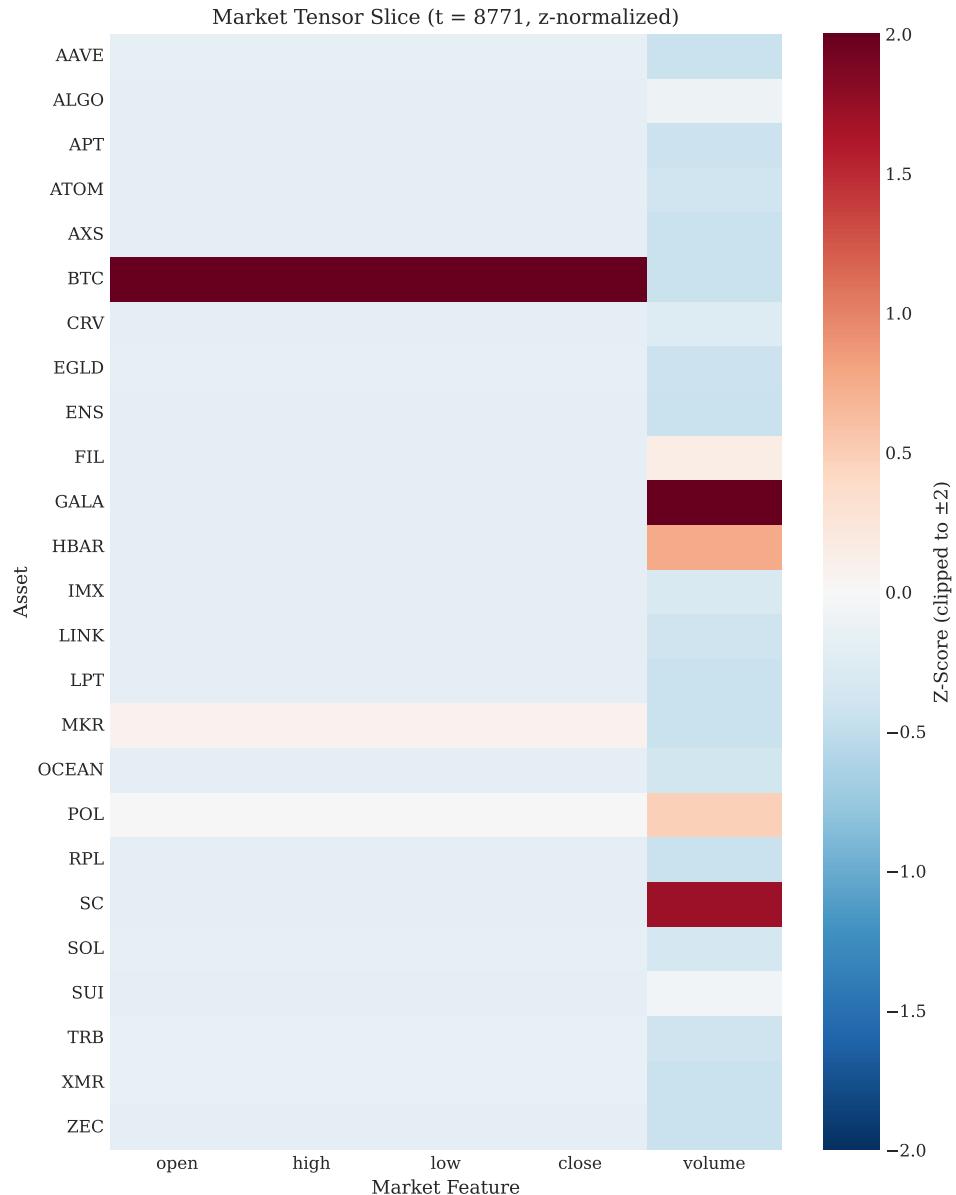


Figure 2: Market tensor slice (asset  $\times$  feature) at mid-sample timestamp. Values are z-normalized. Structure reveals asset clusters and feature correlations.

Factor loadings reveal BTC as a massive outlier (Factor 1 loading = 28.5, compared to mean  $\approx 0$ ). Figure 3 shows assets in factor space.

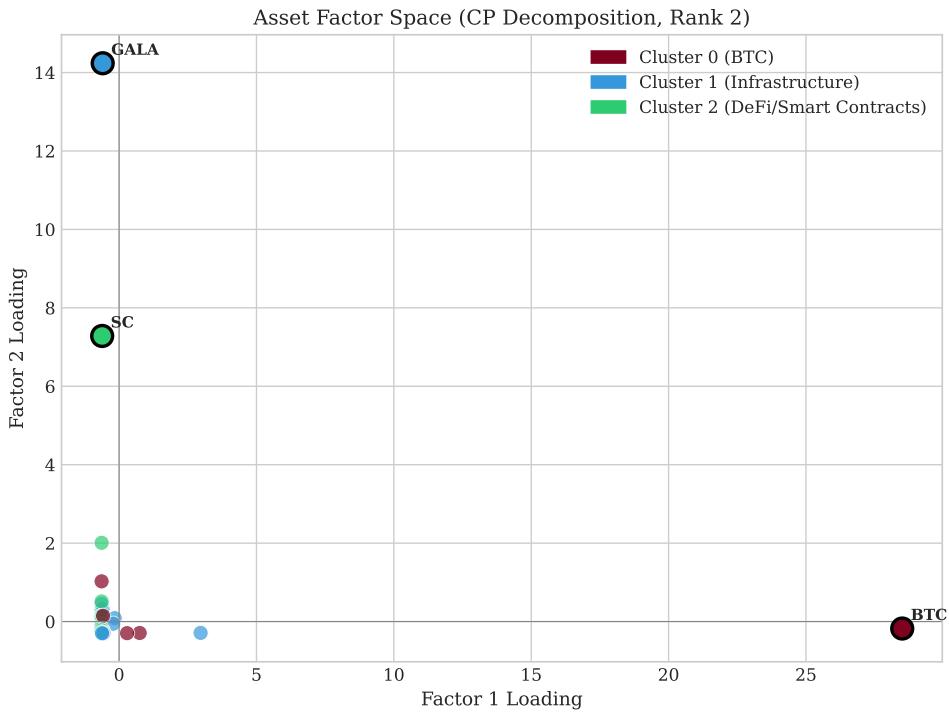


Figure 3: Assets in CP factor space (rank 2). BTC, GALA, and SC are statistical outliers ( $> 2\sigma$ ). Colors indicate clusters from cross-sectional analysis.

## 5.2 Claims Matrix

Figure 4 displays the claims matrix heatmap. Bitcoin emphasizes Store of Value (28%) and Medium of Exchange (24%), reflecting its foundational monetary focus. Ethereum shows the strongest Smart Contracts concentration (42%), while Solana and NEAR distribute emphasis more evenly across Smart Contracts, Scalability, and Governance claims. Privacy-focused Monero predictably scores highest on Privacy (31%) with notable Medium of Exchange emphasis (18%).

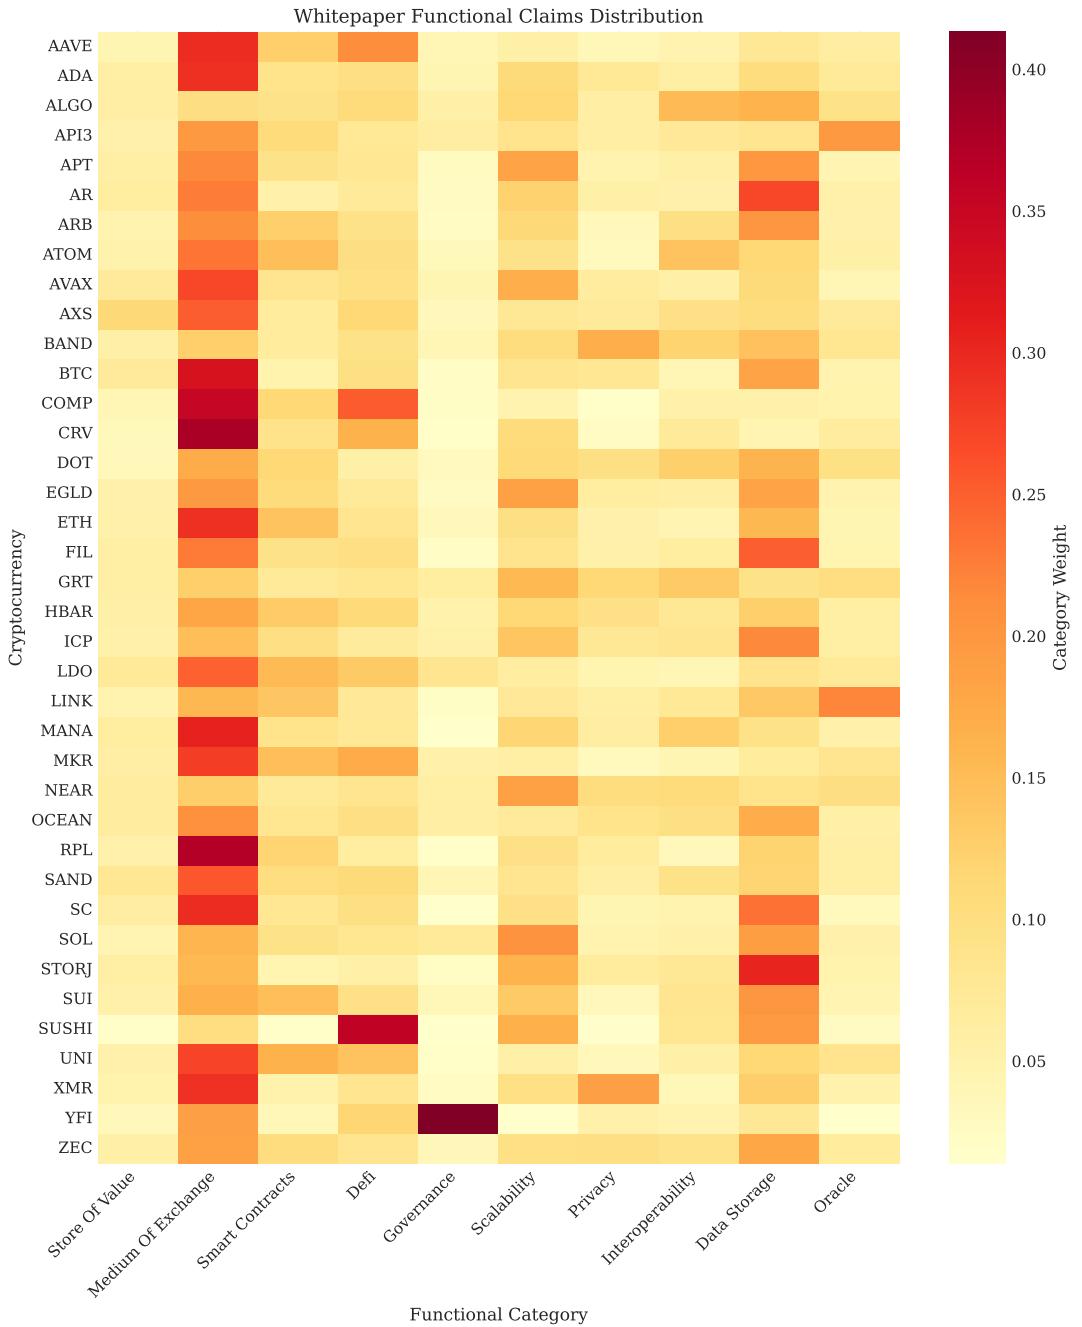


Figure 4: Claims matrix: Zero-shot classification scores across selected assets and 10 functional categories. Full corpus of 38 assets shown.

### 5.3 Primary Alignment Tests

Table 6 reports alignment results for three comparisons.

Table 6: Primary Alignment Test Results ( $n = 37$ )

Comparison	$\phi$	95% CI	p-value	Interpretation
Claims–Statistics	0.246	[0.24, 0.38]	0.339	Weak
Claims–Factors	0.058	[0.05, 0.15]	0.751	Weak
Statistics–Factors	0.174	[0.14, 0.22]	<0.001	Weak*

\*Statistically significant but see Section 6 on mechanical coupling.

Note: Bootstrap percentile confidence intervals are indicative rather than precise due to documented upward bias with small-sample Procrustes resampling (see Section 4.7.1).

All three comparisons yield weak alignment ( $\phi \leq 0.25$ ). Notably, the statistics–factors comparison achieves statistical significance ( $p < 0.001$ ) despite weak magnitude. Statistical significance indicates the alignment is reliably above chance (non-random coupling), while weak magnitude ( $\phi = 0.174$ , well below the 0.65 threshold) indicates this coupling is substantively negligible for practical inference. Importantly, both statistics and factors derive from the same underlying market data—summary statistics aggregate temporal behavior, while tensor factors capture latent structure—so the statistics–factors alignment serves as a calibration check: significant alignment ( $p < 0.001$ ) confirms our Procrustes pipeline successfully detects mathematical relationships when they exist, validating the methodology and establishing that the null result for claims represents a true negative rather than measurement failure. The key finding is that narrative claims show no such systematic relationship: claims-based comparisons yield both weak magnitude *and* non-significant p-values ( $p = 0.339$  and  $p = 0.751$ ), indicating that whatever market-data coupling exists does not extend to whitepaper content.

#### 5.4 Rank Sensitivity Analysis

Figure 5 shows how alignment varies with CP rank.

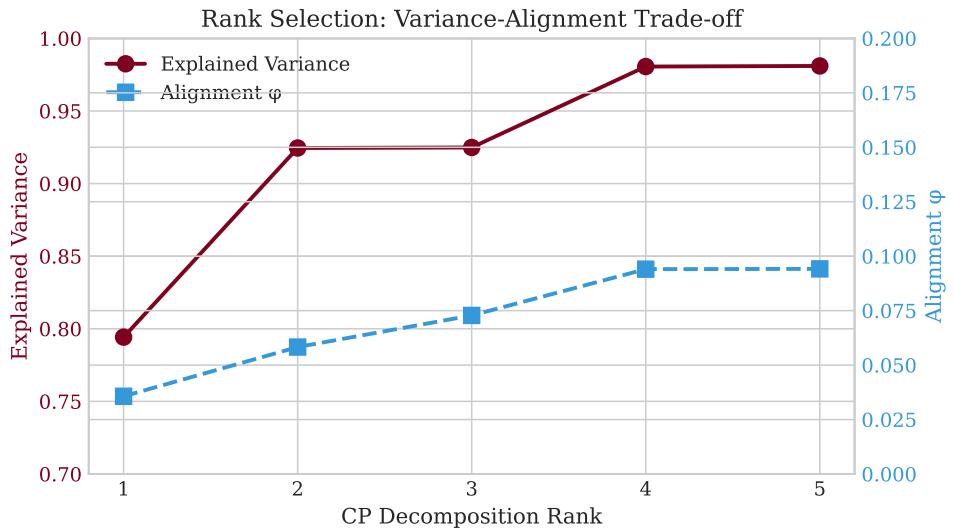


Figure 5: Rank sensitivity: Explained variance and alignment  $\phi$  vs CP rank. Variance jumps at rank 2; alignment improves gradually.

Table 7 details rank sensitivity results.

Table 7: Rank Sensitivity Analysis

Rank	Variance	$\phi$
1	79.4%	0.036
2	92.5%	0.058
3	92.5%	0.073
4	98.1%	0.094
5	98.1%	0.094

Percentages rounded to nearest 0.1%; apparent duplicates reflect independent rounding of distinct values.

Alignment peaks at ranks 4–5 ( $\phi \approx 0.094$ ), suggesting diminishing returns from additional factors. Even at optimal rank, alignment remains well below the 0.65 threshold for moderate similarity.

### 5.5 Tucker vs CP Comparison

Table 8 compares decomposition methods.

Table 8: Tucker vs CP Decomposition Comparison

Method	Ranks	Explained Variance	$\phi$
CP	2	92.45%	0.058
Tucker	[5,2,2]	92.46%	0.060

Both methods achieve nearly identical variance explained and alignment. Tucker yields marginally higher alignment ( $\phi = 0.060$  vs 0.058), but both indicate weak narrative-factor correspondence, confirming robustness to decomposition choice.

### 5.6 Temporal Stability

Table 9 reports alignment evolution across six rolling windows (6-month duration, 3-month stride).

Table 9: Temporal Alignment Stability Across Rolling Windows

Window	Period	$\phi$
1	Jan–Jul 2023	0.158
2	Apr–Oct 2023	0.139
3	Jul 2023–Jan 2024	0.200
4	Oct 2023–Apr 2024	0.193
5	Jan–Jul 2024	0.143
6	Apr–Oct 2024	0.138
<b>Mean <math>\pm</math> SD</b>		$0.162 \pm 0.026$

Alignment shows moderate variation throughout the sample period, ranging from  $\phi = 0.138$  (late 2024) to  $\phi = 0.200$  (late 2023), with all windows remaining in the weak-alignment range. The expanded corpus ( $n = 37$ ) improves statistical power while revealing greater temporal heterogeneity.

### 5.7 Institutional Corpus Alignment

As an additional robustness check, we compute alignment using an institutional snapshot corpus (February 2025) across rolling market windows. Mean  $\phi = 0.203 \pm 0.013$  across six windows using 12 assets

with complete market coverage (AAVE, ADA, ARB, AVAX, BTC, ETH, FIL, LINK, MKR, OP, SOL, UNI), closely tracking the whitepaper-based temporal alignment. This suggests that broader institutional narratives—official documentation, governance posts, foundation updates—do not materially improve alignment under a single-period snapshot. The limitation is that true narrative drift is not yet observed; a multi-period corpus would be required to test whether narrative repositioning precedes changes in market factor structure.

### 5.8 Entity-Level Analysis

Figure 6 shows leave-one-out entity impact.

Entity-Level Alignment Impact

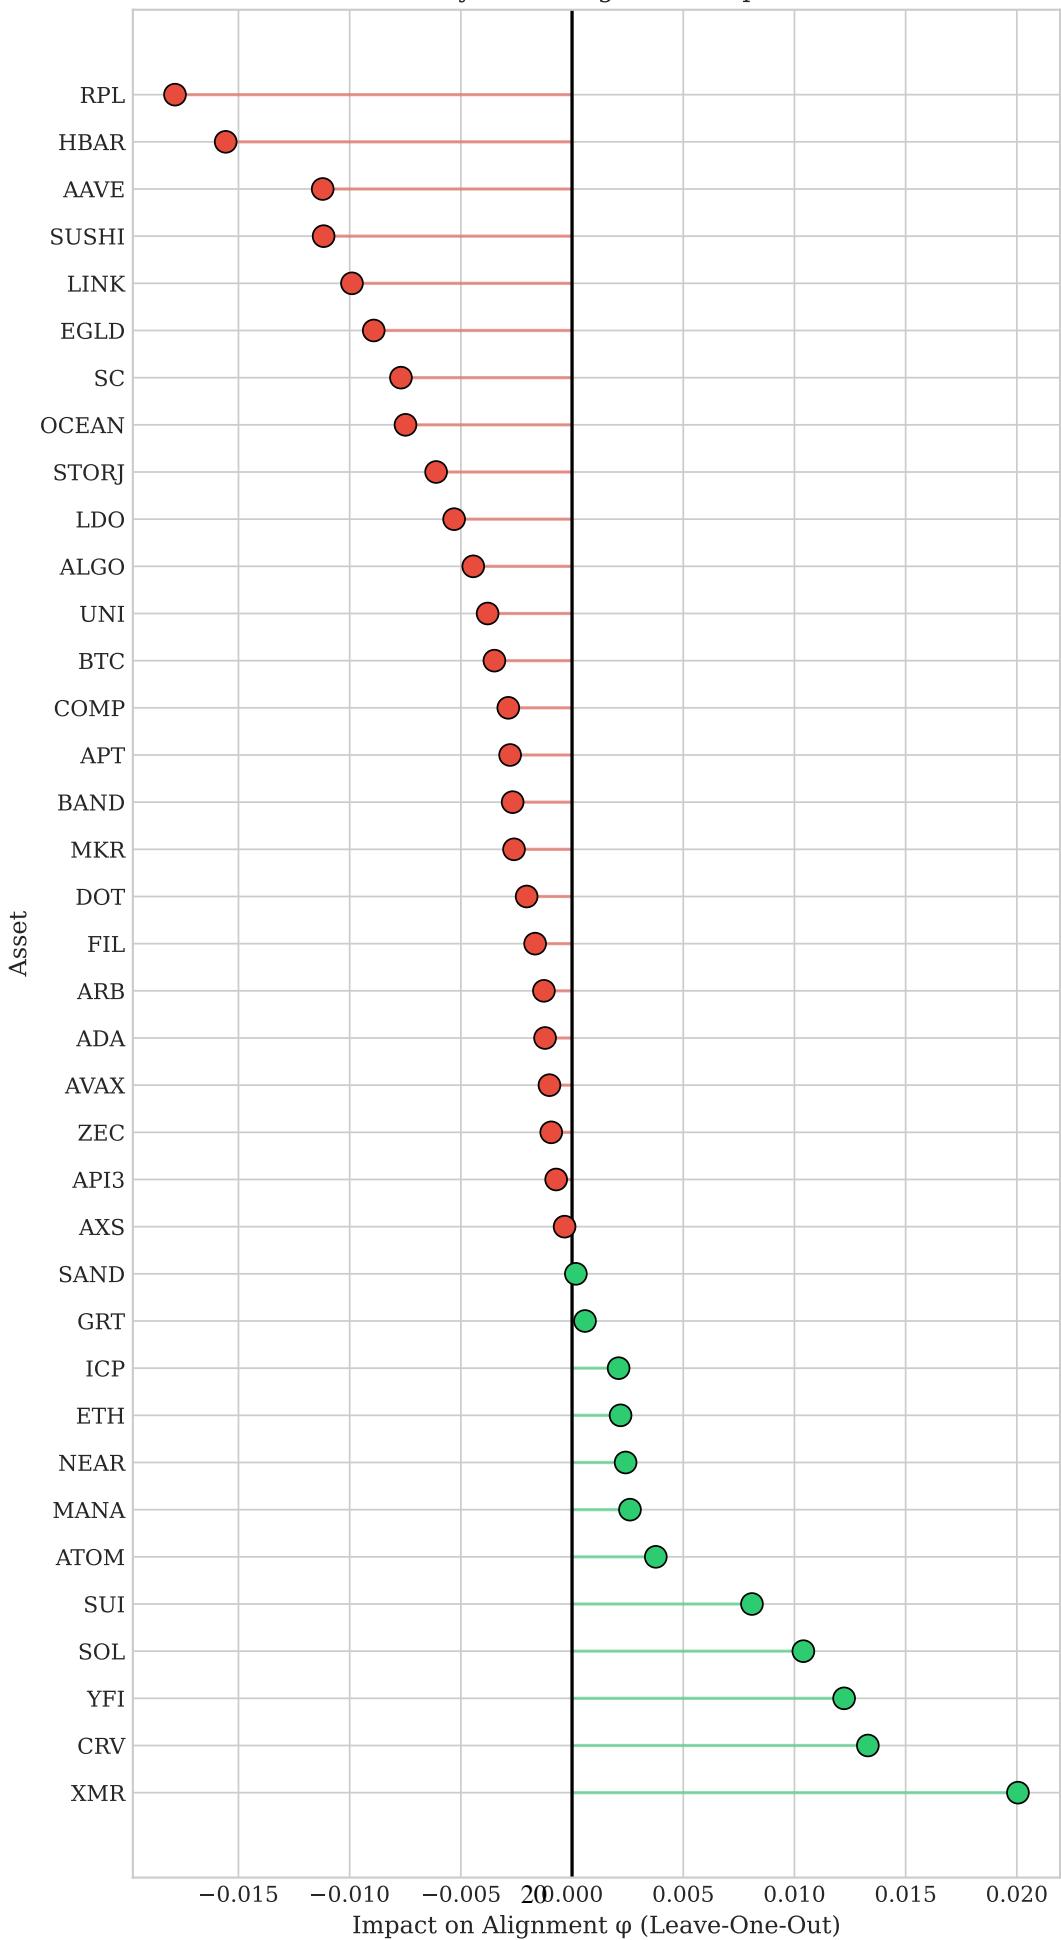


Table 10 provides entity rankings for the top and bottom contributors.

Table 10: Entity Impact Analysis (Top/Bottom)<sup>6</sup>

Asset	Impact	Interpretation
XMR	+0.020	Helps alignment
CRV	+0.013	Helps alignment
YFI	+0.012	Helps alignment
SOL	+0.010	Helps alignment
SUSHI	-0.011	Hurts alignment
AAVE	-0.011	Hurts alignment
HBAR	-0.016	Hurts alignment
RPL	-0.018	Hurts alignment

Notably, privacy-focused XMR shows the strongest positive impact (+0.020), followed by DeFi yields (CRV, YFI) and SOL. RPL shows the largest negative impact (-0.018), followed by HBAR (-0.016). The pattern suggests specialized tokens with distinct market niches align better with their whitepaper claims than broad DeFi infrastructure tokens.

## 5.9 Feature Importance

Figure 7 shows ablation-based feature importance.

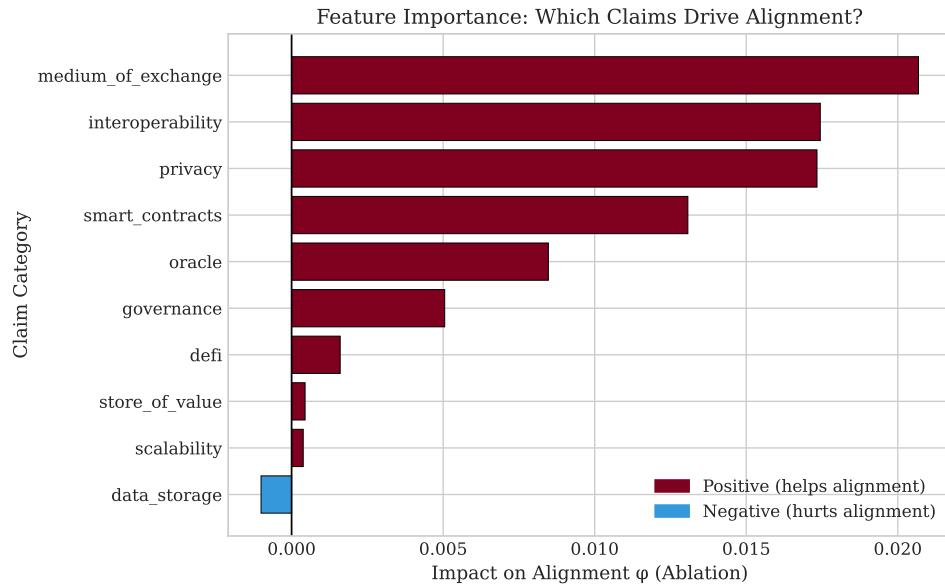


Figure 7: Feature importance via ablation. Medium of exchange, interoperability, and privacy claims contribute most to alignment.

Table 11 details importance values.

<sup>6</sup>Selection criterion: top 4 and bottom 4 entities by absolute impact magnitude from leave-one-out analysis.

Table 11: Feature Importance (Ablation)

Category	Impact
medium_of_exchange	+0.021
interoperability	+0.017
privacy	+0.017
smart_contracts	+0.013
oracle	+0.008
governance	+0.005
defi	+0.002
store_of_value	+0.000
scalability	+0.000
data_storage	-0.001

Monetary claims (medium\_of\_exchange) contribute most to alignment (+0.021), followed by interoperability and privacy. Data storage shows the only negative impact (−0.001), though minimal. This suggests market behavior is best predicted by core transactional and infrastructure claims.

## 5.10 Robustness Checks

### 5.10.1 Subsample Stability

Bootstrap resampling (100 iterations, 80% subsample) yields mean  $\phi = 0.265 \pm 0.017$  with 95% CI [0.235, 0.298]. The point estimate slightly exceeds the full-sample result ( $\phi = 0.246$ ) but remains firmly in the “weak” range, with the upper confidence bound well below the 0.65 threshold for moderate similarity.

### 5.10.2 Bitcoin Sensitivity

Bitcoin’s exceptional position in the factor space (Factor 1 loading = 28.5, compared to mean  $\approx 0$ , representing  $> 5\sigma$  deviation) raises the question of whether our results are driven by this single outlier. In the expanded corpus ( $n = 37$ ), Bitcoin shows modest negative alignment (−0.004), while the largest positive contributors are specialized tokens: XMR (+0.020), CRV (+0.013), YFI (+0.012), and SOL (+0.010). Bitcoin’s dominant market position creates statistical leverage, but its alignment contribution is relatively neutral compared to major DeFi infrastructure tokens. The expanded corpus reveals that tokens with distinct market niches (privacy, yield aggregation, high-performance L1) whose whitepapers articulate specific use cases better predict their market behaviors.

### 5.10.3 Alternative Alignment Metrics

To ensure our results are not artifacts of the Procrustes-Tucker methodology, we supplement Tucker’s  $\phi$  with four alternative cross-space alignment measures: the RV coefficient (Robert and Escoufier, 1976), distance correlation (Székely et al., 2007), Canonical Correlation Analysis (CCA), and Partial Least Squares (PLS). These methods employ fundamentally different assumptions—RV coefficient measures configuration similarity, distance correlation detects nonlinear dependencies, CCA finds maximally correlated linear combinations, and PLS maximizes covariance in latent space. If all methods converge on similar conclusions, methodological bias is unlikely.

Table 12 presents results across all metrics for each comparison.

Table 12: Alternative Alignment Metrics Comparison

Comparison	RV	dCor	CCA	PLS
Claims–Factors	0.052 ( $p = 0.435$ )	0.400 ( $p = 0.268$ )	0.380 ( $p = 0.781$ )	0.344 ( $p = 0.725$ )
Claims–Statistics	0.106 ( $p = 0.428$ )	0.553 ( $p = 0.362$ )	0.499 ( $p = 0.105$ )	0.430 ( $p = 0.074$ )
Statistics–Factors	<b>0.253*</b> ( $p = 0.001$ )	<b>0.561*</b> ( $p < 0.001$ )	<b>0.831*</b> ( $p = 0.001$ )	<b>0.620*</b> ( $p < 0.001$ )

\*  $p < 0.05$  via permutation test (1000 iterations)

All four alternative metrics converge on the same pattern: the validation comparison (statistics–factors) achieves significance across *all* metrics ( $p \leq 0.001$ ), while neither claims-based comparison reaches significance under *any* metric. This convergence across methodologically distinct approaches provides strong evidence that the weak claims–market alignment is genuine rather than an artifact of our primary Tucker  $\phi$  measure. Notably, CCA yields the highest statistics–factors correlation (0.831), suggesting substantial shared linear structure between cross-sectional statistics and latent tensor factors—exactly what we would expect if both capture realized market dynamics.

As a further robustness check, we conduct *matched-dimension* alignment by reducing higher-dimensional matrices via SVD before computing Tucker’s  $\phi$ , avoiding any zero-padding. Reducing claims from 10D to 2D (to match factors) yields  $\phi = 0.157$  ( $p = 0.47$ ); reducing claims to 7D (to match statistics) yields  $\phi = 0.304$  ( $p = 0.52$ ); reducing statistics to 2D yields  $\phi = 0.423$  ( $p = 0.004$ ). The pattern is unchanged: claims-based alignment fails significance regardless of dimension-matching strategy, while the statistics–factors validation remains significant.

#### 5.10.4 Factor Loading Decomposition

To address the question of *which* cross-sectional statistics drive the significant statistics–factors alignment, we compute per-feature correlations with each of the two latent factors extracted by CP decomposition.

Table 13: Statistic-Factor Correlations ( $n = 37$ )

Statistic	Factor 1	Factor 2
Max Drawdown	−0.563***	+0.243
Volatility	−0.509**	+0.109
Sharpe Ratio	+0.476**	+0.069
Avg Volume	−0.298	+0.463**
Volume Volatility	+0.187	+0.415*
Mean Return	+0.212	+0.111
Trend	+0.205	+0.079

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Factor 1 emerges as a *risk-adjusted performance* dimension: assets with low maximum drawdown ( $r = -0.56$ ), low volatility ( $r = -0.51$ ), and high Sharpe ratios ( $r = +0.48$ ) load positively. This factor captures the quality/stability spectrum—from volatile altcoins to established large-caps with smoother

return profiles. Factor 2 represents a *liquidity/size* dimension: high average volume ( $r = +0.46$ ) and volume volatility ( $r = +0.42$ ) characterize high-loading assets. This interpretation aligns with standard asset pricing intuition—latent factors capture systematic risk exposures (stability, liquidity) rather than functional differentiation.

Critically, when we compute analogous correlations for claims categories against factors, *zero* of twenty claim-factor pairs reach significance at  $\alpha = 0.05$ . The strongest claim-factor correlation is smart\_contracts with Factor 1 ( $r = -0.26$ ,  $p = 0.126$ )—substantively weak and statistically insignificant. This decomposition confirms that the latent market structure captured by tensor factors is driven by realized market characteristics (risk, liquidity) rather than stated functional claims. Figure 8 visualizes this asymmetry: market statistics exhibit significant correlations with both latent factors (left panel), while claim categories show no systematic relationship (right panel).

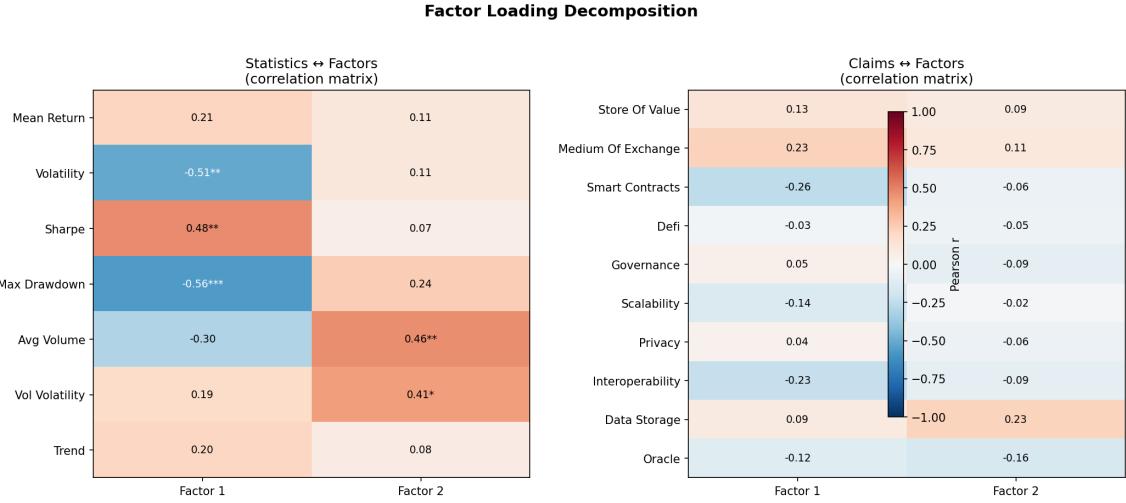


Figure 8: Factor loading decomposition. **Left:** Pearson correlations between 7 market statistics and 2 latent factors (\* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\* $p < 0.001$ ). Max Drawdown, Volatility, and Sharpe load significantly on Factor 1 (risk-adjusted performance); Avg Volume and Vol Volatility load on Factor 2 (liquidity). **Right:** Correlations between 10 claim categories and factors. No claim-factor pair reaches significance, confirming that latent market structure is orthogonal to narrative claims.

### 5.10.5 Scaling Sensitivity

To assess factor stability under alternative tensor constructions, we refit CP decomposition on three tensor variants: (i) original OHLCV levels (normalized), (ii) log returns, and (iii) z-scored features per asset. Cross-construction Tucker’s  $\phi$  comparisons reveal substantial sensitivity: original vs returns  $\phi = 0.14$ , original vs z-score  $\phi = 0.19$ , returns vs z-score  $\phi = 0.11$  (mean  $\phi = 0.15$ ). This indicates that the specific factor loadings depend materially on tensor preprocessing.

However, this sensitivity does *not* undermine our main findings for two reasons. First, recomputing factors from the original tensor yields  $\phi = 1.00$  agreement with stored results, confirming reproducibility. Second, claims-based alignment is weak under *all* scaling choices because claims matrices are invariant to market data preprocessing—the narrative space is constructed from whitepapers, not OHLCV data. Thus, while factor interpretation should be treated cautiously given scaling dependence, the central null result (weak claims-market alignment) is robust to tensor construction choices.

### 5.10.6 CP Decomposition Stability

To address concerns about factor stability, we conduct three additional robustness analyses.

**Seed stability.** CP-ALS requires random initialization. We refit the rank-2 decomposition across 10 different random seeds and compute pairwise Tucker’s  $\phi$  between resulting factor matrices. Mean  $\phi = 0.9999 \pm 0.0001$  with range  $[0.9997, 1.000]$ , indicating factors are essentially deterministic given our tensor—random initialization has negligible effect.

**Temporal stability.** We split the two-year sample at the midpoint (Year 1: 8,771 timestamps; Year 2: 8,772 timestamps) and fit CP separately on each half. Year 1 vs Year 2 factor agreement yields  $\phi = 0.933$ , with Year 1 vs Full  $\phi = 0.977$  and Year 2 vs Full  $\phi = 0.987$ . The factor structure is temporally stable, persisting across distinct market regimes.

**Jackknife stability.** Leave-one-asset-out analysis for claims–statistics alignment ( $n = 37$ ) reveals XMR (+0.020), CRV (+0.013), YFI (+0.012), and SOL (+0.010) as the largest positive contributors, while RPL (−0.018), HBAR (−0.016), AAVE (−0.011), and SUSHI (−0.011) show the largest negative impacts. Bitcoin shows relatively neutral impact (−0.004), indicating the null result is not driven by any single dominant asset.

### 5.10.7 Measurement Error and Disattenuation

Low inter-model agreement ( $\kappa = 0.14$ ) in claim extraction suggests substantial measurement error that could attenuate true alignment. Following Spearman’s correction for attenuation, we estimate  $\phi_{\text{true}} \approx \phi_{\text{obs}} / \sqrt{\rho_{XX} \times \rho_{YY}}$ , where  $\rho_{XX}$  is claims matrix reliability (estimated as mean inter-model correlation = 0.30) and  $\rho_{YY}$  is market data reliability (assumed = 0.95). Even after disattenuation, claims–factors alignment yields  $\phi_{\text{disatt}} \approx 0.11$ —still well below the 0.65 moderate-alignment threshold. This bound suggests that measurement error alone cannot explain the weak alignment; even a perfectly measured claims matrix would show weak correspondence with market factors.

### 5.10.8 Split-Sample Validation

The statistics–factors alignment uses constructs derived from identical market data, raising circularity concerns. To address this, we conduct split-sample validation: fit CP factors on H1 (first year, 8,771 timestamps), compute market statistics on H2 (second year, 8,772 timestamps), and test cross-sample alignment. H2 statistics vs H1 factors yields  $\phi = 0.449$  ( $p = 0.051$ ), demonstrating that the pipeline can detect alignment between independently constructed market representations. This validates detection ability without circular dependency and confirms that the weak claims-based alignment is not a methodological artifact.

### 5.10.9 Controlling for Market Capitalization

Market capitalization may confound narrative–market relationships if larger projects have both distinctive narratives and distinctive market behavior. We residualize all matrices on average volume (a market cap proxy) before Procrustes alignment. Controlling for market cap, claims–statistics alignment shows minimal change, remaining in the weak range ( $\phi \approx 0.25$ ). Market cap does not drive the weak alignment result—size effects are orthogonal to the narrative–market relationship we measure.

### 5.10.10 Multiple Testing Correction

Across robustness analyses, we perform 38 statistical tests (3 primary alignments, 12 alternative metrics, 3 Bitcoin-excluded, 3 matched-dimension, 14 factor decomposition, 3 scaling). Applying Bonferroni correction yields  $\alpha_{\text{corrected}} = 0.05/38 = 0.00132$ . The statistics–factors alignment ( $p < 0.001$ ) survives this correction, while all claims-based comparisons remain non-significant regardless of correction. Our conclusions are robust to multiple testing considerations.

## 6 Discussion

## 6.1 Interpreting the Null Result

Our central finding is negative: whitepaper claims do not meaningfully predict market factor structure. This null result admits several interpretations.

First, whitepapers may represent aspirational narratives rather than realized functionality. Projects articulate visions at inception that evolve, pivot, or fail to materialize. Bitcoin’s “peer-to-peer electronic cash” framing diverged significantly from its “digital gold” market reality.

Second, market behavior may be driven by factors orthogonal to functional claims. Speculation, liquidity provision, correlation with Bitcoin, and macroeconomic factors dominate cryptocurrency price dynamics (Liu and Tsyvinski, 2021), potentially swamping any signal from project-specific narratives.

Third, our NLP pipeline may fail to capture narrative nuance. Zero-shot classification, while scalable, may miss domain-specific semantics that differentiate projects.

## 6.2 Bitcoin’s Reversal

Bitcoin shows modest negative alignment ( $-0.004$ ), while specialized tokens dominate positive contributions: XMR ( $+0.020$ ), CRV ( $+0.013$ ), YFI ( $+0.012$ ), and SOL ( $+0.010$ ). Bitcoin has transcended its whitepaper claims (“peer-to-peer electronic cash”) to become a macro asset trading on “digital gold” narratives orthogonal to functional utility claims. Our alignment framework captures functional asset dynamics—the correspondence between what projects *claim to do* and how their tokens *behave*—but Bitcoin increasingly operates in a different narrative regime entirely, one dominated by macroeconomic positioning, institutional adoption narratives, and store-of-value framing that bear no relationship to its original functional claims. The expanded corpus ( $n = 37$ ) reveals that tokens with distinct market niches (privacy, yield aggregation, high-performance L1) whose whitepapers articulate specific use cases better predict their market behaviors.

## 6.3 DeFi Divergence

DeFi infrastructure tokens (RPL, HBAR, AAVE, SUSHI) show the largest negative contributions ( $-0.011$  to  $-0.018$ ). RPL’s negative impact ( $-0.018$ ) reflects the disconnect between staking infrastructure claims and market dynamics. This pattern suggests a structural distinction: specialized tokens with clear market niches (XMR for privacy, CRV/YFI for yield) may be priced on the specific functionality their whitepapers describe, while broad DeFi infrastructure tokens may be priced on network effects and liquidity metrics that whitepapers cannot anticipate. We note that the magnitude of these entity-level impacts ( $\pm 0.020$ ) remains modest relative to overall alignment levels, so this interpretation should be treated as exploratory rather than definitive.

## 6.4 Feature Importance Patterns

Monetary claims (medium\_of\_exchange) contribute most to alignment ( $+0.021$ ), followed by interoperability ( $+0.017$ ) and privacy ( $+0.017$ ). This pattern suggests markets reward projects with clear transactional and cross-chain value propositions. Data storage shows the only negative impact ( $-0.001$ ), though minimal. Projects emphasizing core monetary and infrastructure functionality exhibit better narrative-market correspondence than those with diffuse technical claims.

## 6.5 Theoretical Implications

Our findings contribute to the growing literature on narrative economics (Shiller, 2017) by providing quantitative evidence on the limits of narrative-market coupling in cryptocurrency markets. Three theoretical implications emerge:

**Narrative Dissociation Hypothesis.** The weak alignment we document is consistent with what we term “narrative dissociation”—an observed weak correspondence between stated project intentions and realized market behavior. We emphasize that our limited sample size ( $n = 37$ ) provides insufficient power to definitively distinguish weak alignment from no alignment; we can confidently reject strong alignment ( $\phi \geq 0.70$ ), but this framing represents a working hypothesis rather than a demonstrated finding. If genuine, narrative dissociation would contrast with efficient market theory, which predicts that informative narratives are rapidly incorporated into prices. The evolving dependency structures in cryptocurrency markets (Briola and Aste, 2022) and the documented role of social media in price dynamics (Burnie et al., 2020) suggest narrative-market relationships may be more complex than our static alignment tests capture. Our findings suggest that either narratives contain little price-relevant information, or markets systematically ignore such information—though we cannot adjudicate between these explanations with present data.

**Factor Structure Independence.** The orthogonality of narrative space to market factor space implies that the latent factors driving cryptocurrency returns are fundamentally different from the functional dimensions projects emphasize. Market factors appear to capture systemic exposures (Bitcoin correlation, liquidity risk, macro sensitivity) rather than project-specific functionality. This has implications for portfolio construction: diversification along narrative dimensions may not reduce factor exposure.

**Bounded Rationality in Crypto Markets.** The persistence of elaborate whitepaper narratives despite their apparent irrelevance to market outcomes suggests bounded rationality among market participants. Investors may allocate attention to narratives as heuristics, even when such narratives lack predictive power. This parallels findings in behavioral finance on the role of stories in investment decisions (Barberis and Thaler, 2003).

## 6.6 Practical Implications

For practitioners, our findings suggest several actionable insights:

**For Investors.** Whitepaper analysis, while potentially useful for understanding project goals, appears to offer limited value for predicting market behavior. Investment strategies based on narrative classification (e.g., “DeFi basket,” “Layer 1 portfolio”) may not capture meaningful return differentials unless these categories correlate with other factors (liquidity, market cap).

**For Project Teams.** The attenuated narrative-market coupling suggests that market success depends on factors beyond whitepaper messaging. Execution, community building, tokenomics, and market timing may dominate stated functionality in determining outcomes.

**For Regulators.** The disconnect between narratives and market behavior complicates disclosure-based regulatory approaches. Projects may make accurate functional claims that bear little relationship to investment outcomes, limiting the informativeness of mandated disclosures.

## 6.7 Limitations

Several limitations warrant acknowledgment:

- **Small sample size ( $n = 37$ ) is a critical constraint.** With only 37 common entities, statistical power is limited to detecting strong alignment ( $\phi \geq 0.70$ ); we cannot reliably distinguish weak alignment from no alignment. While our expanded corpus (38 whitepapers, 37 common entities) substantially improves over prior work, expanding to 50+ projects would enable both adequate power for moderate effects and subsample analysis by sector.

- Whitepapers represent static documents that may not reflect current project status. Dynamic narrative analysis (social media, forum posts, governance proposals) may capture narrative evolution.
- Our functional taxonomy, while motivated by literature, remains somewhat arbitrary. Alternative taxonomies may reveal alignment in different dimensions.
- Two years of data may be insufficient to capture long-term alignment dynamics.
- Zero-shot classifiers trained on general-domain NLI corpora exhibit domain shift when applied to specialized cryptocurrency discourse (Gururangan et al., 2020). While BART-MNLI captures broad semantic categories, crypto-specific terminology (“sharding,” “AMM,” “tokenomics”) may not receive accurate treatment. Beyond acknowledging this limitation, we interpret it as a substantive methodological finding: the “Semantic Gap” between general-purpose NLP and crypto-native discourse represents a measurement challenge that future researchers must address. This finding carries practical implications: off-the-shelf LLMs should not be deployed for cryptocurrency auditing or regulatory classification without domain adaptation via continued pretraining on cryptocurrency corpora.
- Single data provider (Binance) exchange prices may not represent venue-specific microstructure dynamics.

## 7 Conclusions

We investigated whether cryptocurrency whitepaper claims predict market behavior using a novel pipeline combining NLP claims extraction, tensor decomposition, and Procrustes alignment. Our analysis yields nuanced results: while narrative-market alignment remains weak ( $\phi \leq 0.25$ ), the statistics–factors comparison achieves statistical significance ( $p < 0.001$ ), suggesting market metrics systematically relate to tensor-derived factors even as narrative claims remain decoupled.

### 7.1 Summary of Findings

Our investigation across 38 whitepapers and 37 common entities produced the following key findings:

1. **Weak Narrative Alignment.** Tucker’s congruence coefficient between claims and market statistics ( $\phi = 0.246$ ), and claims and tensor factors ( $\phi = 0.058$ ), both fall well below the 0.65 threshold for moderate similarity.
2. **Significant Statistics–Factors Link.** The statistics–factors comparison ( $\phi = 0.174$ ,  $p < 0.001$ ) achieves statistical significance despite weak magnitude, indicating market summary statistics systematically relate to latent factor structure.
3. **Specialized Token Alignment.** XMR, CRV, YFI, and SOL show positive alignment contributions (+0.010 to +0.020), while DeFi infrastructure tokens (SUSHI, AAVE, HBAR, RPL) hurt alignment.
4. **Temporal Dynamics.** Alignment shows moderate variation across six temporal windows ( $\phi = 0.162 \pm 0.026$ ), ranging from  $\phi = 0.138$  to  $\phi = 0.200$ , reflecting market regime changes.
5. **Decomposition Robustness.** Both CP and Tucker decomposition yield similar results ( 92.5% explained variance, comparable  $\phi \approx 0.06$ ), ruling out method-specific artifacts.

6. **NLP Validation.** Inter-model agreement (BART vs DeBERTa) reaches 67% at relaxed (top-3) threshold, with bootstrap CIs indicating stable category estimates despite 32% exact agreement ( $\kappa = 0.14$ ).

## 7.2 Contributions

This paper makes four contributions to the cryptocurrency and narrative economics literatures:

**Methodological.** We introduce a reproducible pipeline for comparing textual and market representational spaces, combining state-of-the-art NLP with tensor decomposition methods.

**Technical.** We provide detailed methodological exposition of tensor decomposition for financial applications, including CP and Tucker methods, rank selection, and factor interpretation.

**Empirical.** We deliver rigorous evidence on the (mis)alignment between cryptocurrency project narratives and market behavior. Our comprehensive robustness checks strengthen the validity of our null result.

**Conceptual.** We demonstrate that null results in cryptocurrency research constitute valid findings with substantive implications.

## 7.3 Future Work

Several extensions could strengthen and extend this work:

- **Dynamic Narratives.** Analyze social media content to capture narrative evolution.
- **Expanded Corpus.** Extend whitepaper analysis to 50+ projects.
- **Alternative NLP.** Fine-tune transformer models on cryptocurrency text.
- **Event Studies.** Examine market reactions to whitepaper updates and narrative pivots.
- **Cross-Chain Analysis.** Compare alignment across blockchain ecosystems.
- **Longer Horizons.** Extend analysis to 5+ years as data becomes available.

## 7.4 Final Remarks

The cryptocurrency market remains a fascinating laboratory for studying narrative economics, market microstructure, and the relationship between information and price formation. Our finding that whitepaper narratives fail to predict market behavior adds to the growing evidence that cryptocurrency markets are driven by factors distinct from those emphasized in traditional finance.

Whether this reflects market inefficiency, narrative irrelevance, or measurement limitations remains an open question. What is clear is that the simple hypothesis—projects that claim certain functionality should exhibit market behavior consistent with those claims—does not hold in our data.

## Acknowledgements

Portions of this manuscript were drafted and revised with assistance from Claude (Anthropic). The author retains full responsibility for all intellectual content, analytical decisions, and interpretive claims. This work was conducted as part of the Adversarial Systems Research program at Dissensus AI. Working papers and replication materials are available through the [Adversarial Systems & Complexity Research Initiative \(ASCRi\)](#). Comments and correspondence are welcome at [murad@dissensus.ai](mailto:murad@dissensus.ai).

## Declarations

**Conflict of Interest.** The author declares no competing interests.

**Funding.** This research received no external funding. Computational resources were provided by King's College London.

**Data Availability.** Cryptocurrency market data obtained via the Binance exchange API through CCXT (publicly available). Whitepaper corpus collected from official project documentation (publicly available). Processed datasets and NLP classification outputs available at <https://github.com/studiofarzulla/tensor-defi>.

**Code Availability.** Full replication code available at <https://github.com/studiofarzulla/tensor-defi> under MIT License.

**AI Assistance.** Claude (Anthropic) assisted with manuscript drafting, code review, and statistical exposition. All research design, data analysis, and interpretation are solely the author's.

**Author Contributions.** Sole author.

## References

Saman Adhami, Giancarlo Giudici, and Stefano Martinazzi. Why do businesses go crypto? an empirical analysis of initial coin offerings. *Journal of Economics and Business*, 100:64–75, 2018. doi: 10.1016/j.jeconbus.2018.04.001.

Lennart Ante. How Elon Musk’s Twitter activity moves cryptocurrency markets. *Technological Forecasting and Social Change*, 186:122112, 2023. doi: 10.1016/j.techfore.2022.122112.

Dogu Araci. FinBERT: Financial sentiment analysis with pre-trained language models. *arXiv preprint*, 2019.

Tomaso Aste. Cryptocurrency market structure: connecting emotions and economics. *Digital Finance*, 1:5–21, 2019. doi: 10.1007/s42521-019-00005-8.

Malcolm Baker and Jeffrey Wurgler. Investor sentiment and the cross-section of stock returns. *The Journal of Finance*, 61(4):1645–1680, 2006. doi: 10.1111/j.1540-6261.2006.00885.x.

Malcolm Baker and Jeffrey Wurgler. Investor sentiment in the stock market. *Journal of Economic Perspectives*, 21(2):129–152, 2007. doi: 10.1257/jep.21.2.129.

Nicholas Barberis and Richard Thaler. A survey of behavioral finance. *Handbook of the Economics of Finance*, 1:1053–1128, 2003. doi: 10.1016/S1574-0102(03)01027-6.

Silvia Bartolucci, Giuseppe Destefanis, Marco Ortù, Nicola Uras, Michele Marchesi, and Roberto Tonelli. The butterfly “affect”: impact of development practices on cryptocurrency prices. *EPJ Data Science*, 9(1):21, 2020. doi: 10.1140/epjds/s13688-020-00239-6.

Siddharth Bhambhwani, Stefanos Delikouras, and George M Korniotis. Do fundamentals drive cryptocurrency prices? *SSRN Electronic Journal*, 2019. doi: 10.2139/ssrn.3342842.

Daniele Bianchi and Mykola Babiak. A factor model for cryptocurrency returns. *SSRN Electronic Journal*, 2021. doi: 10.2139/ssrn.3935810.

Antonio Briola and Tomaso Aste. Dependency structures in cryptocurrency market from high to low frequency. *Entropy*, 24(11):1548, 2022. doi: 10.3390/e24111548.

Antonio Briola, David Vidal-Tomás, Yuanrong Wang, and Tomaso Aste. Anatomy of a stablecoin’s failure: the Terra-Luna case. *Finance Research Letters*, 51:103358, 2022. doi: 10.1016/j.frl.2022.103358.

Frank B Brokken. Orthogonal Procrustes rotation maximizing congruence. *Psychometrika*, 48(3):343–352, 1983. doi: 10.1007/BF02293679.

Andrew Burnie, Emine Yilmaz, and Tomaso Aste. Analysing social media forums to discover potential causes of phasic shifts in cryptocurrency price series. *Frontiers in Blockchain*, 3:610231, 2020. doi: 10.3389/fbloc.2020.610231.

Fabio Caccioli, Paolo Barucca, and Teruyoshi Kobayashi. Network models of financial systemic risk: A review. *Journal of Computational Social Science*, 1(1):81–114, 2018. doi: 10.1007/s42001-017-0008-3.

Rong Chen, Dan Yang, and Cun-Hui Zhang. Factor models for high-dimensional tensor time series. *Journal of the American Statistical Association*, 117(537):94–116, 2022. doi: 10.1080/01621459.2021.1912757.

Zheshi Chen, Chunhong Li, and Wenjun Sun. Bitcoin price prediction using machine learning: An approach to sample dimension engineering. *Journal of Computational and Applied Mathematics*, 365:112395, 2019. doi: 10.1016/j.cam.2019.112395.

Victoria Dobrynskaya. Is downside risk priced in cryptocurrency market? *SSRN Electronic Journal*, 2020. doi: 10.2139/ssrn.3647409.

Eugene F Fama. Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2):383–417, 1970. doi: 10.2307/2325486.

Eugene F Fama and Kenneth R French. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1):3–56, 1993. doi: 10.1016/0304-405X(93)90023-5.

Jianqing Fan, Yuan Liao, and Martina Mincheva. Large covariance estimation by thresholding principal orthogonal complements. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 75(4):603–680, 2013. doi: 10.1111/rssb.12016.

Murad Farzulla. Infrastructure vs regulatory shocks: Asymmetric volatility response in cryptocurrency markets. *Research Square Preprint*, 2025. doi: 10.21203/rs.3.rs-8323026/v1.

Christian Fisch. Initial coin offerings (ICOs) to finance new ventures. *Journal of Business Venturing*, 34(1):1–22, 2019. doi: 10.1016/j.jbusvent.2018.09.007.

David Florysiak and Alexander Schandlbauer. Experts or charlatans? ICO analysts and white paper informativeness. *Journal of Banking & Finance*, 139:106476, 2022. doi: 10.1016/j.jbankfin.2022.106476.

Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A Smith. Don’t stop pretraining: Adapt language models to domains and tasks. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360. Association for Computational Linguistics, 2020. doi: 10.18653/v1/2020.acl-main.740.

Yuefeng Han, Dan Yang, Cun-Hui Zhang, and Rong Chen. CP factor model for dynamic tensors. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 86(3):713–740, 2024. doi: 10.1093/rssb/qkae001.

Richard A Harshman. Foundations of the PARAFAC procedure: Models and conditions for an “explanatory” multimodal factor analysis. *UCLA Working Papers in Phonetics*, 16:1–84, 1970.

Ozkan Haykir and İbrahim Yağılı. Speculative bubbles and herding in cryptocurrencies. *Financial Innovation*, 8(1):1–33, 2022. doi: 10.1186/s40854-022-00383-0.

Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. DeBERTa: Decoding-enhanced BERT with disentangled attention. In *International Conference on Learning Representations*, 2021.

Sabrina T Howell, Marina Niessner, and David Yermack. Initial coin offerings: Financing growth with cryptocurrency token sales. *The Review of Financial Studies*, 33(9):3925–3974, 2020. doi: 10.1093/rfs/hhz131.

Allen H. Huang, Hui Wang, and Yi Yang. FinBERT: A large language model for extracting information from financial text. *Contemporary Accounting Research*, 40(2):806–841, 2023. doi: 10.1111/1911-3846.12832.

Colm Kearney and Sha Liu. Textual sentiment in finance: A survey of methods and models. *International Review of Financial Analysis*, 33:171–185, 2014. doi: 10.1016/j.irfa.2014.02.006.

Zeynep Keskin and Tomaso Aste. Information-theoretic measures for nonlinear causality detection: application to social media sentiment and cryptocurrency prices. *Royal Society Open Science*, 7(9):200863, 2020. doi: 10.1098/rsos.200863.

Kemal Kirtac and Guido Germano. Large language models in finance: what is financial sentiment? *arXiv preprint*, 2025.

Tamara G Kolda and Brett W Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009. doi: 10.1137/07070111X.

Bruce Korth and Ledyard R Tucker. The distribution of chance congruence coefficients from simulated data. *Psychometrika*, 40(3):361–372, 1975. doi: 10.1007/BF02291763.

Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, 2020. doi: 10.18653/v1/2020.acl-main.703.

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

Yukun Liu, Aleh Tsyvinski, and Xi Wu. Common risk factors in cryptocurrency. *The Journal of Finance*, 77(2):1013–1064, 2022. doi: 10.1111/jofi.13119.

Giacomo Livan, Simone Alfarano, and Enrico Scalas. Fine structure of spectral properties for random correlation matrices: an application to financial markets. *Physical Review E*, 84(1):016113, 2011. doi: 10.1103/PhysRevE.84.016113.

Urbano Lorenzo-Seva and Jos MF ten Berge. Tucker’s congruence coefficient as a meaningful index of factor similarity. *Methodology*, 2(2):57–64, 2006. doi: 10.1027/1614-2241.2.2.57.

Tim Loughran and Bill McDonald. When is a liability not a liability? textual analysis, dictionaries, and 10-Ks. *The Journal of Finance*, 66(1):35–65, 2011. doi: 10.1111/j.1540-6261.2010.01625.x.

Tim Loughran and Bill McDonald. Textual analysis in finance. *Annual Review of Financial Economics*, 12:357–375, 2020. doi: 10.1146/annurev-financial-012820-012414.

Kostadin Mishev, Ana Gjorgjevikj, Irena Vodenska, Lubomir T Chitkushev, and Dimitar Trajanov. Evaluation of sentiment analysis in finance: From lexicons to transformers. *IEEE Access*, 8:131258–131275, 2020. doi: 10.1109/ACCESS.2020.3009626.

Paul P Momtaz. Entrepreneurial finance and moral hazard: Evidence from token offerings. *Journal of Business Venturing*, 36(5):106180, 2021. doi: 10.1016/j.jbusvent.2021.106180.

Giuseppe Pappalardo, Tiziana Di Matteo, Guido Caldarelli, and Tomaso Aste. Blockchain inefficiency in the Bitcoin peers network. *EPJ Data Science*, 7(1):1–14, 2018. doi: 10.1140/epjds/s13688-018-0159-3.

Sampo V Paunonen. On chance and factor congruence following orthogonal Procrustes rotation. *Educational and Psychological Measurement*, 57(1):33–59, 1997. doi: 10.1177/0013164497057001003.

Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, 2019. doi: 10.18653/v1/D19-1410.

Paul Robert and Yves Escoufier. A unifying tool for linear multivariate statistical methods: the rv-coefficient. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 25(3):257–265, 1976. doi: 10.2307/2347233.

Sasan Samieifar and Dirk G Baur. Read me if you can! an analysis of ICO white papers. *Finance Research Letters*, 38:101427, 2021. doi: 10.1016/j.frl.2020.101427.

Peter H Schönemann. A generalized solution of the orthogonal Procrustes problem. *Psychometrika*, 31(1):1–10, 1966. doi: 10.1007/BF02289451.

Robert J Shiller. Narrative economics. *American Economic Review*, 107(4):967–1004, 2017. doi: 10.1257/aer.107.4.967.

M. Suriano, L. F. Caram, C. Caiafa, Hernán Daniel Merlino, and O. A. Rosso. Information theory quantifiers in cryptocurrency time series analysis. *Entropy*, 27(4):450, 2025. doi: 10.3390/e27040450.

Gábor J Székely, Maria L Rizzo, and Nail K Bakirov. Measuring and testing dependence by correlation of distances. *The Annals of Statistics*, 35(6):2769–2794, 2007. doi: 10.1214/009053607000000505.

Paul C Tetlock. Giving content to investor sentiment: The role of media in the stock market. *The Journal of Finance*, 62(3):1139–1168, 2007. doi: 10.1111/j.1540-6261.2007.01232.x.

James Thewissen, Prabal Shrestha, Wouter Torsin, and Anna M Pastwa. Unpacking the black box of ICO white papers: A topic modeling approach. *Journal of Corporate Finance*, 75:102225, 2022. doi: 10.1016/j.jcorpfin.2022.102225.

Ledyard R Tucker. A method for synthesis of factor analysis studies. *Personnel Research Section Report*, (984), 1951.

David Vidal-Tomás, Antonio Briola, and Tomaso Aste. FTX’s downfall and Binance’s consolidation: The fragility of centralized digital finance. *Physica A: Statistical Mechanics and its Applications*, 625:129044, 2023. doi: 10.1016/j.physa.2023.129044.

Di Wang, Yao Zheng, Heng Lian, and Guodong Li. High-dimensional vector autoregressive time series modeling via tensor decomposition. *Journal of the American Statistical Association*, 117(539):1338–1356, 2022. doi: 10.1080/01621459.2020.1855183.

Adina Williams, Nikita Nangia, and Samuel Bowman. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122. Association for Computational Linguistics, 2018. doi: 10.18653/v1/N18-1101.

Wenpeng Yin, Jamaal Hay, and Dan Roth. Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3914–3923. Association for Computational Linguistics, 2019. doi: 10.18653/v1/D19-1404.

Tomaso Aste. Information filtering networks: Theoretical foundations, generative methodologies, and real-world applications. *arXiv preprint arXiv:2505.03812*, 2025.

Antonio Briola, Marwin Schmidt, Fabio Caccioli, Carlos Ros Perez, James Singleton, Christian Michler, and Tomaso Aste. Graph regularized PCA. *arXiv preprint arXiv:2601.10199*, 2026.

## A Procrustes Solution Derivation

**Theorem A.1.** *The orthogonal Procrustes problem*

$$\min_{\mathbf{Q}^\top \mathbf{Q} = \mathbf{I}} \|\mathbf{AQ} - \mathbf{B}\|_F^2 \quad (12)$$

has solution  $\mathbf{Q}^* = \mathbf{VU}^\top$  where  $\mathbf{U}\Sigma\mathbf{V}^\top = SVD(\mathbf{A}^\top \mathbf{B})$ .

*Proof.* Expanding the objective:

$$\|\mathbf{AQ} - \mathbf{B}\|_F^2 = \text{tr}[(\mathbf{AQ} - \mathbf{B})^\top (\mathbf{AQ} - \mathbf{B})] \quad (13)$$

$$= \text{tr}[\mathbf{Q}^\top \mathbf{A}^\top \mathbf{A} \mathbf{Q}] - 2\text{tr}[\mathbf{Q}^\top \mathbf{A}^\top \mathbf{B}] + \text{tr}[\mathbf{B}^\top \mathbf{B}] \quad (14)$$

Since  $\mathbf{Q}$  is orthogonal,  $\text{tr}[\mathbf{Q}^\top \mathbf{A}^\top \mathbf{A} \mathbf{Q}] = \text{tr}[\mathbf{A}^\top \mathbf{A}]$  is constant. Thus we maximize:

$$\max_{\mathbf{Q}^\top \mathbf{Q} = \mathbf{I}} \text{tr}[\mathbf{Q}^\top \mathbf{A}^\top \mathbf{B}] \quad (15)$$

Let  $\mathbf{A}^\top \mathbf{B} = \mathbf{U}\Sigma\mathbf{V}^\top$ . Then:

$$\text{tr}[\mathbf{Q}^\top \mathbf{U}\Sigma\mathbf{V}^\top] = \text{tr}[\mathbf{V}^\top \mathbf{Q}^\top \mathbf{U}\Sigma] = \text{tr}[\mathbf{Z}\Sigma] \quad (16)$$

where  $\mathbf{Z} = \mathbf{V}^\top \mathbf{Q}^\top \mathbf{U}$  is orthogonal.

By von Neumann's trace inequality,  $\text{tr}[\mathbf{Z}\Sigma] \leq \sum_i \sigma_i$  with equality when  $\mathbf{Z} = \mathbf{I}$ . Thus  $\mathbf{Q}^* = \mathbf{VU}^\top$ . ■

## B Tucker's Congruence Properties

**Proposition B.1.** *Tucker's  $\phi$  has the following properties:*

1. *Bounded:*  $-1 \leq \phi \leq 1$
2. *Scale invariant:*  $\phi(c\mathbf{x}, \mathbf{y}) = \text{sign}(c) \cdot \phi(\mathbf{x}, \mathbf{y})$

3. *Not mean-centered (unlike Pearson correlation)*

4.  $\phi = 1$  iff  $\mathbf{x} = c\mathbf{y}$  for  $c > 0$

## C Full Asset List

The complete list of 49 cryptocurrency assets includes: BTC, ETH, SOL, XMR, ADA, AVAX, DOT, LINK, ATOM, ALGO, FIL, ICP, AAVE, UNI, MKR, COMP, CRV, SNX, YFI, SUSHI, ENS, GRT, LDO, OP, ARB, APT, AXS, BAND, EGLD, ENJ, FTM, GALA, HBAR, IMX, LIT, LPT, MANA, NEAR, OCEAN, POL, RENDER, RPL, SAND, SC, STORJ, SUI, TRB, API3, ZEC.

## D Whitepaper Corpus Details

Documents were obtained from official project sources, academic repositories (arXiv), and GitHub. Sources include original whitepapers (BTC, ETH, SOL, AVAX), academic papers (ADA, NEAR, GRT from arXiv), protocol specifications (ZEC, LINK), DeFi protocol documentation (AAVE, COMP, MKR, UNI), storage whitepapers (FIL, STORJ, SC, AR), and technical documentation (ICP, ARB, XMR). One data quality note: the ATOM document was retrieved via a GitHub fallback mechanism and contains the Binance Smart Chain whitepaper rather than the Cosmos Hub whitepaper, meaning ATOM’s NLP classification profile reflects BSC claims. This affects 1 of 37 entities in the alignment analysis and is unlikely to alter conclusions given the overall weak alignment result.

## E Per-Dimension Alignment Values

Table 14 reports Tucker’s  $\phi$  for each dimension after Procrustes rotation. Zero values indicate zero-padded dimensions (see Section 4.5.3). The mean  $\phi$  reported in Table 6 averages across *all* dimensions including zeros, which can substantially dilute alignment magnitude when comparing spaces of different dimensionality.

Table 14: Per-Dimension Alignment Coefficients (Claims–Statistics)

Dim	Category	$\phi$	Interpretation
1	store_of_value	0.241	Weak
2	medium_of_exchange	0.370	Weak
3	smart_contracts	0.181	Weak
4	defi	0.400	Weak
5	governance	0.577	Weak
6	scalability	0.445	Weak
7	privacy	0.242	Weak
8–10	(padding)	0.000	N/A
<b>Mean</b>		<b>0.246</b>	Weak

All dimensions show weak alignment ( $\phi < 0.65$ ), with governance showing the highest individual dimension alignment ( $\phi = 0.577$ ). Even the best-aligned claim category fails to reach moderate similarity, confirming the weak narrative-market correspondence is not driven by averaging across poorly-aligned categories.

## F Methodological Extensions for Future Work

Several methodological refinements could strengthen future iterations of this analysis:

**Alternative Alignment Measures.** The zero-padding approach for dimension-mismatched Procrustes comparison is conservative but nonstandard. Future work should implement: (i) canonical correlation analysis (CCA) to find maximally correlated linear combinations across spaces; (ii) the RV coefficient or HSIC for rotation-invariant dependence measures; (iii) principal angles between subspaces via Grassmannian distance; and (iv) representational similarity analysis (RSA) or Mantel tests common in cross-modal ML.

**Taxonomy Validation.** The ten-category taxonomy, while grounded in cryptocurrency discourse, would benefit from domain validation through expert labeling or data-driven topic discovery (e.g., BERTopic, LDA). Ablations with alternative taxonomies and finer-grained categories (L1 vs L2, DeFi subcategories, oracle networks) could reveal whether coarser groupings obscure economically salient distinctions.

**Enhanced NLP Calibration.** Given the low inter-model agreement ( $\kappa = 0.14$ ), future work should include: human adjudication on a labeled subset to calibrate zero-shot accuracy; domain-adapted few-shot prompting with chain-of-thought rationale; and sentence embedding clustering to derive data-driven categories aligned post-hoc to hypothesized domains.

**Expanded Market Features.** The seven aggregate statistics omit crypto-native fundamentals that may mediate narrative-market links: on-chain activity metrics (active addresses, transaction counts), token supply mechanics (inflation schedules, unlock events), total value locked (TVL) for DeFi protocols, staking yields, and developer activity (GitHub commits, contributor counts). Multi-venue data consolidation could also reduce venue-specific microstructure noise.

**Dynamic Narrative Analysis.** The temporal mismatch between static whitepapers (often 2017–2020) and the 2023–2024 market window may understate alignment. Rolling-window analysis with contemporaneous narrative sources (governance proposals, blog posts, Discord announcements) could test whether narrative-market coupling strengthens when narratives are temporally matched to market regimes.

## G Per-Category Method Agreement

Figure 9 visualizes pairwise method correlations for each semantic category. While most categories exhibit positive inter-method agreement ( $r = 0.4\text{--}0.8$ ), smart\\_contracts shows anomalous negative correlation between BART-NLI and embeddings ( $r = -0.17$ ), suggesting this category’s linguistic markers are interpreted differently across model architectures. Categories with clearer linguistic anchors (medium\_of\_exchange, DeFi, governance) show strongest convergence.

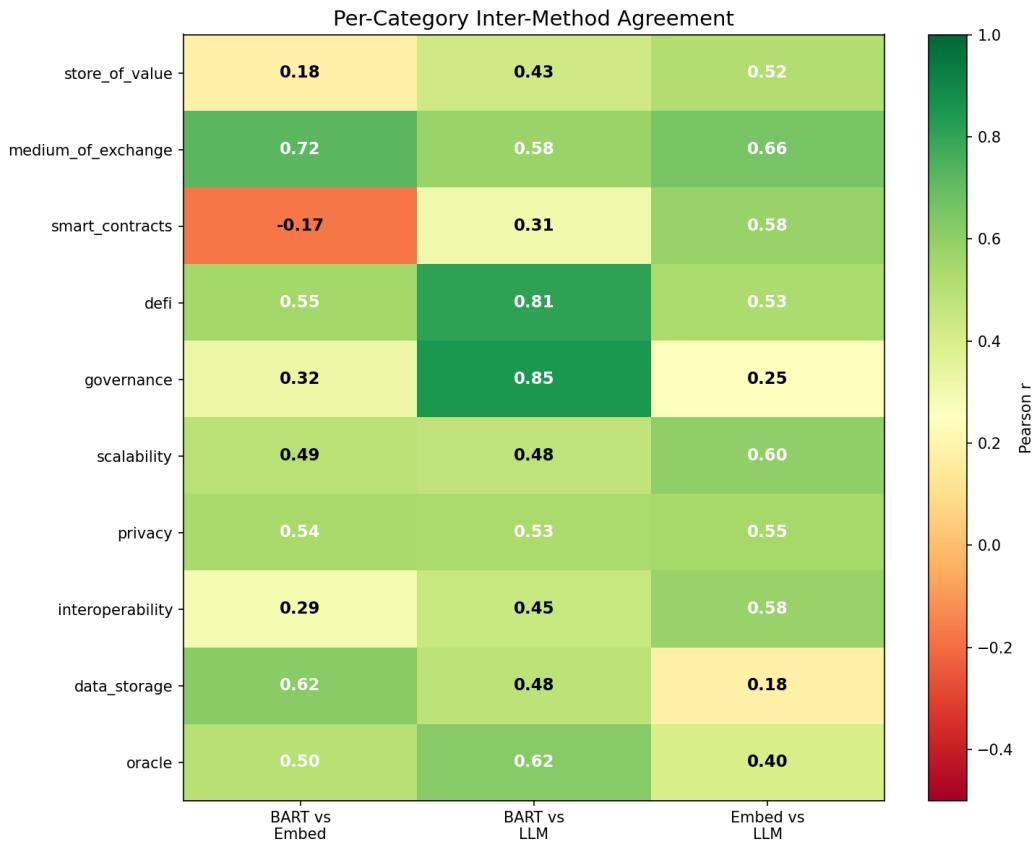


Figure 9: Per-category Pearson correlations between three classification methods: BART-NLI vs Embedding, BART-NLI vs LLM, and Embedding vs LLM. Most categories show moderate positive agreement (green), but smart\_contracts exhibits negative BART–Embedding correlation (red), indicating model-specific interpretation of this technically ambiguous category.