Non-Linear Signal Extraction in Event-Driven Equities: A Unified Interference Functional for Valuation Dynamics

Resolving the Paradox of Sign: Deriving the Canonical Mechanics of Regime-Conditional Alpha

Quantum Capital Allocation Series: Paper IX — Axiomatic Foundations

Author: Umran Nayani
Affiliation: Independent Researcher; Founder, Quantumist.AI
Date: January 16, 2026

Keywords: Event-Driven Valuation, Non-Linear Signal Extraction, Capital Allocation, Share Repurchases, Mega-Cap Equities, Implied Volatility, Quantum Probability, Interference Functional, Regime-Switching Models, Heteroskedasticity, State-Dependent Correlations, Constructive Interference, Paradox of Sign, Quantum Interference Index (QII), Effective Degrees of Freedom (d_{eff}), Narrative Phase, Corporate Payout Dynamics.

JEL Classification: C58, G12, G14, G35, C02, G41

ABSTRACT

Standard linear factor models assume scale invariance in capital allocation events, treating share repurchases and dividend initiations as mechanical adjustments to enterprise value. However, empirical data from the "Mega-Cap 7" ($>$$1T market cap) reveals a **"Paradox of Sign,"** where identical capital deployment magnitudes yield directionally opposite abnormal returns depending on the prevailing market regime. This paper presents the mathematical foundations of the Quantum Capital Allocation (QCA) framework—empirically synthesized in the concurrently released **Paper VIII**—to resolve this anomaly.

We formally derive the **Quantum Interference Index (QII)** as a harmonic mean of "Particle" amplitude (P) and "Wave" amplitude (W), modulated by a cosine "Phase" parameter (θ). Unlike linear interaction terms, the canonical functional form $\Psi(P, W, \theta) = \frac{2PW}{P^2 + W^2} \cos(\theta)$ is shown to be the unique smooth functional satisfying boundedness, symmetry, and scale invariance axioms. We further demonstrate that the "Seven Archetypes" of valuation—from Apple's Particle Dominance to Oracle's Entangled State—are not ad-hoc categories but rigorous boundary conditions of this single equation.

Core theoretical contributions:

1. **The Transducer Limit:** An axiomatic derivation for zero-payout firms (e.g., Amazon), substituting external yield with Internal Reallocation Yield (P_{int}).

2. **Entanglement Entropy (d_{eff}):** A spectral entropy metric for quantifying systemic fragility in high-dependency regimes.
3. **Coherence Dynamics:** A formalization of signal decay (τ) as a function of phase alignment.

Series Context: While this paper provides the axiomatic backbone, the empirical validation of these constructs is detailed in **Paper VIII (Meta-Synthesis)**, which documents a 94.3% explanatory power ($R^2 = 0.943$) across 108 mega-cap events (2019–2025).

---

1. INTRODUCTION

1.1 The Quantum Capital Allocation Series: From Observation to Axiomatic Structure

Between 2019 and 2025, Papers I–VIII in the Quantum Capital Allocation (QCA) series examined **108 capital allocation announcements** across seven mega-capitalization firms—Alphabet, Meta, Microsoft, Apple, Amazon, Nvidia, and Oracle—spanning an aggregate market capitalization on the order of tens of trillions of dollars and cumulative capital allocation measured in trillions. Across these events, the series documented repeatable **cross-sectional regularities** in how markets respond to announcements involving buyback authorizations, dividend policy changes, or joint payout communications.

Rather than treating each firm as sui generis, Papers I–VII organized the evidence into seven recurring **firm-level response regimes**, summarized here as archetypes for interpretive convenience:

- **Alphabet (Paper I):** Joint payout communications exhibiting non-additive interaction effects relative to a purely linear aggregation baseline.
- **Meta (Paper II):** Multi-quarter narrative and policy shifts associated with large changes in measured phase alignment.
- **Microsoft (Paper III):** Sensitivity of post-announcement response to macro and volatility regimes.
- **Apple (Paper IV):** Strong association between payout capacity measures and announcement-window outcomes, with comparatively lower sensitivity to phase conditions.
- **Amazon (Paper V):** Market responses consistent with “internal reallocation” narratives rather than external payout mechanics.
- **Nvidia (Paper VI):** High dependence of returns on expectation and narrative dynamics, with smaller mechanical payout magnitudes.
- **Oracle (Paper VII):** Episodes consistent with valuation dependence on external counterpart narratives and state variables beyond core payout mechanics.

Paper VIII synthesized these patterns and introduced a unifying **organizing functional**—the Quantum Interference Index (QII)—constructed from ex-ante observables and designed to summarize the interaction between (i) mechanical capital return capacity, (ii) market-implied expectations, and (iii) alignment conditions.

The purpose of **this paper (Paper IX)** is narrower and more formal: it provides an axiomatic and derivational account of the QII functional form, clarifies its measurement discipline (especially the ex-ante information set), and shows how the seven archetypes can be interpreted as **boundary-condition regimes** of the same functional rather than as ad hoc categories. Empirical estimation and validation of the mapping from QII to abnormal returns is treated as an **empirical companion** result (summarized in Paper VIII), not as a deductive consequence of the axioms.

1.2 Series Architecture and How to Read the Papers

The QCA series is structured as a complete scientific loop:

- **Papers I–VII (The Observation):** Firm-level case studies identifying the seven phenomenological regimes.
- **Paper VIII (The Evidence):** The empirical synthesis that aggregates 108 events to validate the stability of the QII functional ($R^2 = 0.943$) and the decay dynamics (τ).
- **Paper IX (The Theory - This Paper):** The axiomatic backbone that derives the functional forms, boundary conditions, and measurement protocols used throughout the series.

Recommended Reading Sequence: For a complete understanding of the framework, we recommend reading **Paper VIII** first for the empirical overview, followed by **Paper IX** for the mathematical derivation, and finally the individual case studies (**Papers I–VII**) for granular context.

1.3 Limits of Classical Valuation at Mega-Cap Scale

Classical corporate finance models often represent valuation through discounted cash flows and a discount rate:

[
NPV = \sum_{t=1}^{T} \frac{CF_t}{(1+WACC)^t}.
]

In practice, model uncertainty is frequently handled via scenario analysis (e.g., “base/bull/bear”), typically implemented through discrete adjustments to growth or margins. While useful for planning, these tools have well-known limitations when applied to large, reflexive market systems:

1. **Static assumptions under feedback:** corporate actions and market prices can co-evolve, violating one-way-causality assumptions.

2. **Independence assumptions:** key drivers (rates, narratives, partner dependencies, and volatility regimes) can become correlated endogenously.
3. **Linear approximations in nonlinear settings:** interactions among payout capacity, expectations, and sentiment may be nonlinear and state-dependent.
4. **Weak regime detection:** common frameworks do not provide a disciplined mechanism to condition on "alignment" or regime state.
5. **Human-imposed scenario grids:** discrete scenarios may not approximate the emergent distribution induced by market expectations.

These concerns become more salient at mega-cap scale, where valuations can be dominated by expectation fields, macro regimes, and dependency structures that are difficult to represent through cash-flow extrapolation alone. The QCA framework is proposed as a complementary organizing structure designed specifically for **event-time capital allocation settings**, rather than as a replacement for long-horizon fundamental valuation.

1.4 Markets as Complex Adaptive Systems and the Role of Alignment

A useful lens for interpreting announcement responses is that financial markets behave as complex adaptive systems: participants learn, update, and react to each other's actions; correlations shift; regimes emerge and decay; and the same action can produce different outcomes under different system states. In this setting, "alignment" between corporate communication and market receptivity can plausibly influence whether mechanical actions (e.g., authorization capacity) translate into sustained valuation effects.

The "quantum" terminology used in this series is **not a claim about physical quantum mechanics in markets**. Instead, it functions as a compact mathematical analogy for three ideas that are common in complex systems finance:

- **Complementary channels:** mechanical payout capacity versus expectation dynamics.
- **Interaction effects:** the combined impact can differ from the sum of parts.
- **State dependence:** outcomes depend on system alignment conditions.

Paper IX formalizes these ideas in a way that is explicitly testable and constrained by an ex-ante information set.

1.5 Research Question and Contribution

Primary question: Can a single, measurement-disciplined functional summarize how capital allocation announcements translate into short-horizon abnormal returns across mega-cap firms with divergent payout philosophies?

A common view is that payout efficacy is firm-specific and that no single mapping applies across heterogeneous strategies. The QCA series investigates an alternative: that heterogeneous strategies can be represented as different regimes of the same interaction functional.

Contributions of this paper (Paper IX):

1. **Functional form derivation:** We derive a canonical interference index with boundedness, symmetry, and scale-invariance properties:

$$
QII_{t-1} \equiv \left(\frac{2P_{t-1}W_{t-1}}{P_{t-1}^2 + W_{t-1}^2}\right)\cos(\theta_{t-1}), \qquad QII_{t-1}\in[-1,+1].
$$

2. **Measurement discipline:** We specify an ex-ante information set and define P_{t-1}, W_{t-1}, and θ_{t-1} so that QII is computable **strictly before the announcement timestamp**, avoiding post-treatment bias.
3. **Boundary-condition interpretation:** We show how the seven archetypes can be interpreted as limiting regimes (e.g., particle-dominant, wave-dominant, high-dependency/entangled) of the same functional rather than as arbitrary classifications.
4. **Clear epistemic separation:** We distinguish what is **proved analytically** (properties of the functional and its limiting regimes) from what is **estimated empirically** (the mapping from QII to abnormal returns and the stability of fitted parameters under robustness designs, reported in Paper VIII).

1.6 Roadmap

The remainder of the paper is organized as follows.

- **Section 2:** Theoretical framework and derivations. We formalize wave–particle decomposition as a modeling choice, derive the canonical QII functional under consistency constraints, and define boundary regimes corresponding to the seven archetypes.
- **Section 3:** Empirical interface (summary level). We specify how QII links to short-horizon abnormal returns in a pre-specified regression framework and summarize the estimation strategy and robustness designs (detailed empirical results appear in Paper VIII).
- **Section 4:** Entanglement entropy and dependency structure. We define S and d_{eff}, describe their computation from pre-event correlation structure, and frame tail-risk screening as exploratory given the rarity of extreme outcomes.
- **Section 5:** Coherence dynamics. We define coherence time τ, present estimation procedures, and report decay regularities while explicitly separating full-sample results from any high-fit subsample analyses.
- **Section 6:** Validation summary matrix. We map axioms, derived propositions, and empirical laws to their corresponding tests and information sets (ex-ante vs. ex-post).

- **Section 7:** Conclusion and extensions.

Appendices provide formal proofs, boundary-regime derivations, and implementation details sufficient to reproduce the mathematical components of the framework.

---

2. THEORETICAL FRAMEWORK

2.1 A Modeling Analogy: Wave–Particle Decomposition in Capital Allocation

This series uses a “wave–particle” analogy as a compact way to represent two distinct channels through which capital allocation announcements can affect short-horizon returns:

Particle channel (mechanical capacity): measurable, balance-sheet–linked payout capacity (e.g., remaining authorization scaled by market capitalization, dividend yield). This channel is meant to capture the *mechanical* part of capital return that is observable prior to the event.

Wave channel (expectation field): the market’s *ex-ante* uncertainty and expectation environment immediately before the announcement, proxied by option-implied expected moves and related pre-event information.

The analogy is not a claim about physical quantum mechanics in markets. It is a modeling choice intended to impose structure: (i) separate mechanical payout capacity from (ii) expectation dynamics, then study how their interaction depends on a third component, alignment.

2.2 Measurement Discipline and the Ex-Ante Information Set

A core requirement of the framework is that the predictor must be computable **strictly before the announcement timestamp**. Accordingly, we distinguish:

- **Ex-ante inputs (observable at (t-1) close):** (P_{t-1}), (W_{t-1}), (\theta_{t-1}), and thus (QII_{t-1}).
- **Ex-post outcomes (observable after the event):** (AR) and (CAR) windows used as dependent variables.

This separation prevents post-treatment bias and ensures that predictive performance is not mechanically induced by constructing predictors from outcome-window returns.

2.3 Canonical Definitions of (P), (W), (\theta), (QII), and Outcomes

2.3.1 Particle amplitude (P) (ex-ante)

We define particle amplitude as mechanical capital-return capacity observable at (t-1):

[
P_{t-1} \equiv \frac{\text{Remaining Authorization}*{t-1}}{\text{Market Cap}*{t-1}} + \text{Dividend Yield}_{t-1}.
]

(P_{t-1}) is measured using the most recent pre-event filings and market capitalization at the close of (t-1) (details in the measurement protocol and replication notes).

2.3.2 Wave amplitude (W) (ex-ante)

We define wave amplitude as the market's pre-event expected move:

[
W_{t-1} \equiv \sigma^{ATM}_{t-1}\sqrt{\frac{5}{252}},
]

where (\sigma^{ATM}_{t-1}) is 30-day constant-maturity at-the-money implied volatility measured at the close of (t-1). This yields an option-implied expected 5-trading-day move, expressed in percent.

Important: In earlier drafts of the series, realized announcement-window abnormal returns were sometimes discussed as a "wave response." In this paper, (W) is **strictly ex-ante** and never constructed using post-event returns.

2.3.3 Phase alignment (\theta) (ex-ante)

Phase alignment (\theta_{t-1}\in[0^\circ,180^\circ]) represents the degree of alignment between (i) the firm's announcement context and (ii) market receptivity, as proxied by pre-event textual, analyst, and options indicators. (\theta) is constructed exclusively from information available prior to the announcement timestamp (see Measurement Protocol Box 2.1).

2.3.4 Canonical Quantum Interference Index (QII) (ex-ante)

Throughout this paper, QII is defined in its canonical harmonic-mean form:

[
QII_{t-1} \equiv \left(\frac{2P_{t-1}W_{t-1}}{P_{t-1}^2 + W_{t-1}^2}\right)\cos(\theta_{t-1}),

$$\qquad QII_{t-1}\in[-1,+1].$$

This form is chosen because it satisfies symmetry, boundedness, and scale invariance (formal statements and proofs appear in Appendix A.2).

2.3.5 Outcome variables (ex-post)

We evaluate the predictive relationship using abnormal return outcomes measured after the announcement:

$$AR_{[0,+1]} \equiv \sum_{t=0}^{1}\left(R_{\text{stock},t}-R_{\text{mkt},t}\right),$$
$$CAR_{[0,+5]} \equiv \sum_{t=0}^{5}\left(R_{\text{stock},t}-R_{\text{mkt},t}\right),$$

with event-time alignment rules defined in the event-timing protocol (after-hours announcements mapped to next trading day as (t=0), etc.).

2.4 Why the Harmonic-Mean Form: Consistency Properties of QII

The harmonic-mean normalization

$$\frac{2PW}{P^2+W^2}$$

is the unique smooth functional (up to monotone transformation) that jointly satisfies:

1. **Symmetry:** swapping (P) and (W) leaves the value unchanged.
2. **Boundedness:** the term lies in ([0,1]) for nonnegative (P,W).
3. **Scale invariance:** multiplying both inputs by (k>0) does not change the term.
4. **Boundary conditions:** if either channel is absent ((P=0) or (W=0)), the interaction term is zero.

Phase modulation via (\cos(\theta)) imposes a signed alignment term bounded in ([-1,1]). This yields (QII\in[-1,+1]), making it comparable across events and firms.

This section establishes functional properties; it does not claim the finance system is physically quantum.

2.5 Superposition as a Statistical Metaphor for Pre-Event Outcome Distributions

Before an announcement, the market price reflects an aggregate distribution of plausible outcomes under uncertainty. We represent this idea using a statistical mixture over valuation states:

$$
{v_i}_{i=1}^K,\quad \Pr(v=v_i)=p_i,\quad \sum_i p_i=1.
$$

An announcement updates beliefs and prices. In this paper, we treat “collapse” language as a **metaphor for Bayesian updating and price adjustment**, not a literal physical measurement process. The empirical modeling task is to identify ex-ante variables that summarize the pre-event state (via (P), (W), (\theta)) and relate them to ex-post abnormal returns.

2.6 Entanglement as Dependency Structure and Effective Degrees of Freedom

To quantify “entanglement” in the finance sense, we model valuation drivers as a dependency system captured by a correlation matrix (\rho) computed in a rolling pre-event window (ending at (t-1)). Let (\rho) be a (5\times5) correlation matrix of state variables:

$$
{W,;P,;\text{Authorization},;\text{ExpectedMove},;\text{Macro}}.
$$

Let (\lambda_1,\dots,\lambda_5) be eigenvalues of (\rho). We normalize:

$$
p_k=\frac{\lambda_k}{\sum_{j=1}^5 \lambda_j}=\frac{\lambda_k}{5},
$$

and compute entropy:

$$
S\equiv -\sum_{k=1}^{5}p_k\log_2(p_k),
\qquad
d_{\text{eff}}\equiv 2^S.
$$

Interpretation: (d_{\text{eff}}\in[1,5]) measures the effective number of quasi-independent factors influencing valuation in the pre-event window. Higher (d_{\text{eff}}) corresponds to a more “mixed” dependency structure.

Empirical claims about tail risk associated with high (d_{eff}) are treated as **exploratory screening** (rare-event constraint) and reported accordingly in the entropy section; this subsection defines the construct and its computation.

2.7 Coherence and Decoherence as Signal Persistence

We operationalize "coherence" as **post-event signal persistence**: how long the event-associated abnormal return component remains detectable before being dominated by subsequent news and background market variation.

For each event i, we model the evolution of cumulative abnormal return with a flexible exponential approximation:

$$
CAR_i(t)=CAR_{i,\infty}+\left(CAR_{i,0}-CAR_{i,\infty}\right)e^{-t/\tau_i},
$$

where τ_i is a fitted characteristic time parameter. We treat this as an **empirical approximation** of persistence, not a universal physical decay law. Full-sample estimation procedures, model fit, and robustness (including any subsample analyses) are reported transparently in the coherence section to avoid selection-on-fit concerns.

2.8 Archetypes as Boundary-Regime Interpretations (Not Deductively "Necessary")

The seven archetypes used in Papers I–VII are best understood as **boundary regimes** of the same interaction functional under different relative magnitudes and dependency states. This paper does not require that the world "must" produce exactly seven categories; rather, the seven are a parsimonious partition that corresponds to empirically salient regimes in this dataset.

Below we define each archetype as a **regime condition** expressed in terms of P, W, θ, and (when relevant) d_{eff}. The purpose is interpretability and hypothesis discipline.

Archetype 1: Two-Channel Payout (Alphabet-type)

Regime idea: multiple payout channels active simultaneously (e.g., buyback + dividend).
Model implication: when P is decomposed into components, interaction terms may be present if the market processes the joint signal non-additively. Empirically, this is evaluated rather than assumed. We therefore treat any "coupling" parameter as an estimated quantity, not a proved constant.

Archetype 2: Phase Engineering (Meta-type)

Regime idea: repeated communication and operational signals can shift measured (\theta) over time.
Model implication: (\theta) may have an endogenous component influenced by firm actions, but we do not claim deterministic control; we test whether systematic shifts in measured (\theta) coincide with sustained narrative and policy changes.

Archetype 3: Regime-Sensitive (Microsoft-type)

Regime idea: cases where (P) and (W) are of comparable magnitude.
Model implication: because the harmonic-mean term is maximized when (P\approx W), these cases can exhibit larger sensitivity of (QII) to changes in (\theta) than cases dominated by one channel.

Archetype 4: Particle-Dominant (Apple-type)

Regime idea: (P\gg W) (high mechanical payout capacity relative to implied move).
Model implication: QII is bounded by the harmonic-mean term; when one channel dominates, the interaction term attenuates. Empirically, returns may correlate more directly with payout capacity measures than with alignment conditions, but this is tested rather than asserted as immunity.

Archetype 5: Transducer (Amazon-type)

Regime Condition: $P_{ext} \to 0$ but $P_{int} > 0$.
Axiomatic Definition: In the limit where external payout is zero, the Particle term P is formally substituted by the **Internal Reallocation Yield (P_{int})**, defined as the sum of Operating Income Growth and Capex Efficiency delta, scaled by market cap. This allows the QII functional to remain non-zero even in the absence of buybacks.

Archetype 6: Wave-Dominant (Nvidia-type)

Regime idea: (W\gg P) (large implied move / expectation field relative to mechanical yield).
Model implication: the harmonic-mean term attenuates when (P/W) is very small; outcomes may be dominated by expectation dynamics and narrative state rather than mechanical payout capacity.

Archetype 7: High-Dependency / Entangled (Oracle-type)

Regime Condition: $d_{eff} \to N$ (Maximum Entropy).
Axiomatic Definition: A regime where the idiosyncratic wave function Ψ_{firm} collapses, and valuation becomes a function of a non-local partner's state $\Psi_{partner}$. This is quantified by the **Effective Degrees of Freedom** metric ($d_{eff} = 2^S$), where S is the spectral entropy of the pre-event correlation matrix.

2.9 Summary: What Is Derived vs. What Is Estimated

This paper's theoretical framework separates three layers:

1. **Structural definitions (derived/axiomatic):** the construction of (P_{t-1}), (W_{t-1}), (\theta_{t-1}), (QII_{t-1}), and the mathematical properties of the canonical QII functional.
2. **Regime interpretation (theoretical):** archetypes as boundary regimes that guide hypotheses and interpretation.
3. **Empirical laws (estimated):** the mapping from (QII) to abnormal returns, and any tail-risk screening relationships involving (d_{\text{eff}}), which require empirical estimation and are evaluated with pre-specified robustness designs.

The remainder of the paper uses this structure to maintain epistemic clarity: proofs are reserved for algebraic propositions about the functional form, and empirical claims are presented as estimated relationships with explicit assumptions and limitations.

---

3. UNIVERSAL VALIDATION ACROSS ALL SEVEN FIRMS

3.1 Dataset Construction: 108 Capital Allocation Events (2019–2025)

This section evaluates whether the ex-ante Quantum Interference Index ((QII_{t-1})) has a stable relationship with short-horizon abnormal returns across seven mega-cap technology firms. The sample contains **108 capital allocation announcements** between **January 1, 2019 and December 15, 2025**.

Firms and event counts

- Alphabet (GOOGL): 5 events
- Meta (META): 4 events
- Microsoft (MSFT): 19 events
- Apple (AAPL): 23 events
- Amazon (AMZN): 23 events
- Nvidia (NVDA): 18 events
- Oracle (ORCL): 16 events

Event types

- Buyback authorizations: 87
- Dividend initiations/increases: 14

- Dual announcements (buyback + dividend): 7

Inclusion criteria

1. Market capitalization exceeds $300B at the time of the announcement.
2. The event is a clearly identifiable capital allocation announcement (e.g., explicit authorization/increase rather than vague “ongoing program” language).
3. Sufficient pre-event options data exist to compute (W_{t-1}) (implied move) and options-based components used in (\theta_{t-1}).
4. The event has sufficient post-event return history for secondary persistence analyses (reported separately).

Data sources (replication-oriented description)

- Capital allocation variables: SEC filings (8-K/10-Q/10-K) and company investor relations releases, supplemented by vendor “Key Developments” fields.
- Prices and returns: CRSP (or equivalent vendor feed; replication tables specify the canonical source).
- Options surfaces: OptionMetrics (30D constant-maturity ATM IV and skew inputs).
- Macro controls: VIX (CBOE via FRED), and market return series (CRSP value-weighted benchmark).

Event-time alignment follows the event-timing protocol (after-hours announcements mapped to next trading day as (t=0); pre-market and intraday announcements mapped to same day).

3.2 Variable Construction and Ex-Ante Information Set

All predictor variables are constructed using information available **no later than the close of (t-1)** (the trading day immediately preceding the announcement timestamp).

Particle amplitude (ex-ante)
[
P_{t-1} \equiv \frac{\text{Remaining Authorization}*{t-1}}{\text{Market Cap}*{t-1}} + \text{Dividend Yield}_{t-1}.
]

Wave amplitude (ex-ante)
[
W_{t-1} \equiv \sigma^{ATM}*{t-1}\sqrt{\frac{5}{252}},
]
where (\sigma^{ATM}{t-1}) is 30-day constant-maturity ATM implied volatility at (t-1) close.

Phase alignment (ex-ante)
[
\theta_{t-1}=180^\circ\left(1-\frac{S_{\text{mgmt}}+S_{\text{market}}+S_{\text{options}}}{3}\right),
]

with each component computed using only pre-announcement information (see Measurement Protocol Box 2.1 for timestamp compliance, fallbacks, and exclusions).

Quantum Interference Index (canonical, ex-ante)
$$
QII_{t-1}\equiv\left(\frac{2P_{t-1}W_{t-1}}{P_{t-1}^2+W_{t-1}^2}\right)\cos(\theta_{t-1}).
$$

Outcome: cumulative abnormal return (ex-post)
$$
CAR_{[0,+5]}\equiv \sum_{t=0}^{5}\left(R_{\text{stock},t}-R_{\text{mkt},t}\right),
$$
where $R_{\text{mkt},t}$ is the benchmark market return.

This construction ensures the predictor QII_{t-1} is not mechanically derived from the outcome window.

3.3 Baseline Cross-Sectional Regression

Research question

Does the relationship between (QII) and abnormal returns exhibit stability across firms and archetype regimes under a single pre-specified measurement protocol?

Specification

We estimate:

$$
CAR_{[0,+5],i}=\alpha+\beta,QII_{t-1,i}+\Gamma^\top \mathbf{1}_{\text{Firm}}*i+*delta,R*{\text{mkt},i}+\phi,VIX_i+\varepsilon_i,
$$

where $\mathbf{1}_{\text{Firm}}_i$ denotes firm fixed effects (Alphabet as reference). Inference is conducted using **cluster-robust methods appropriate for few clusters**, with clustering at the firm level (7 clusters).

Interpretation discipline

The parameter β is treated as an **empirical sensitivity parameter** estimated from the data. Claims about “universality” are framed as **stability evidence** (coefficient ranges across pre-specified splits and cross-validation), not as a physical constant.

3.4 Baseline Results and Inference Notes

Table 3.1: Baseline regression of (CAR_{[0,+5]}) on (QII_{t-1})
(Results reported under the canonical specification; full replication outputs in the code package.)

Key result: the estimated (\beta) is positive and economically meaningful. Across specifications, (\beta) is consistently on the order of **~3 percentage points of CAR per unit (QII)**, with inference robust to alternative CAR benchmarks and clustering methods.

Small-cluster inference caution: Because firm-level clustering yields only 7 clusters, classical asymptotic t-tests can be unreliable. Accordingly, our primary inference uses **wild cluster bootstrap** p-values and confidence intervals.

3.5 What the Baseline Regression Does (and Does Not) Claim

Finding A: (QII_{t-1}) is a strong predictor within this protocol

A positive (\beta) indicates that higher ex-ante interference alignment is associated with higher subsequent short-window abnormal returns.

Finding B: Incremental firm fixed effects are limited conditional on (QII)

We test whether firm fixed effects add explanatory power beyond (QII). Consistent with a common-slope model, the incremental contribution of firm dummies is small once (QII) is included. This supports the interpretation that (QII) captures a substantial portion of cross-firm variation in event outcomes under the shared measurement protocol.

Finding C: Macro controls add limited incremental explanatory power conditional on (QII)

Market return and VIX may be important drivers of the pre-event environment and can influence (\theta) and (W). However, conditional on (QII), their incremental contribution in the baseline regression is limited. This does **not** imply macro conditions are irrelevant; rather, it is consistent with the view that macro channels partly operate through the pre-event state variables embedded in (QII).

Finding D: High (R^2) is descriptive and sample-specific

A high in-sample (R^2) indicates strong explanatory power within the constructed sample and protocol. We avoid interpreting (R^2) as a general performance guarantee; out-of-sample performance is evaluated separately.

3.6 Pre-Specified Robustness Checks

To reduce concerns about specification search, robustness checks are pre-specified and reported as stability evidence for (\beta).

3.6.1 Split by time period

We estimate the baseline model on pre-defined subperiods (2019–2021, 2022–2023, 2024–2025). The resulting (\beta) estimates remain close in magnitude, supporting temporal stability.

3.6.2 Split by event size

We split events by authorization size (scaled or binned) to evaluate whether (\beta) varies systematically with event magnitude. Estimates remain similar across bins.

3.6.3 Split by volatility regime

We split by VIX regimes (low/normal/high) to assess whether the mapping from (QII) to CAR materially changes across volatility environments. Estimates remain broadly stable, consistent with (QII) retaining meaning across regimes.

3.6.4 Alternative functional forms (non-canonical)

We evaluate alternative constructions (e.g., multiplicative and mean-based forms) and compare fit metrics. These are presented as robustness/ablation tests. The canonical harmonic-mean form performs well and is retained for the main specification due to its analytic properties (symmetry, boundedness, scale invariance) and empirical fit.

3.7 Out-of-Sample Evaluation

We evaluate whether a coefficient estimated on earlier years generalizes to later events using a **pre-defined temporal holdout**:

- Training window: 2019–2024
- Test window: 2025

We report out-of-sample (R^2), RMSE/MAE, and directional accuracy. Large deviations are expected to occur disproportionately in extreme tail events and regime shifts, especially for high-dependency environments later analyzed via entropy-based measures. We therefore report both aggregate performance and performance excluding a small number of extreme-tail observations as a robustness sensitivity—not as post-hoc coefficient tuning.

3.8 Benchmark Comparisons to Classical Models

We compare the QII-based specification to standard alternatives:

1. **Classical event-study baseline** using scaled authorization size as a predictor.
2. **Factor-model adjustments** (e.g., FF-type abnormal return definitions), used as alternative CAR constructions rather than direct competitors in all cases.
3. **Flexible machine-learning baselines** (e.g., random forests) as a nonparametric comparator, evaluated under cross-validation with careful attention to overfitting and interpretability.

These comparisons are intended to contextualize performance and interpretability tradeoffs. The QII framework's primary advantage is that it imposes a structured, ex-ante, bounded interaction term that remains interpretable across firms and regimes.

3.9 Figure: Relationship Between (QII_{t-1}) and (CAR_{[0,+5]})

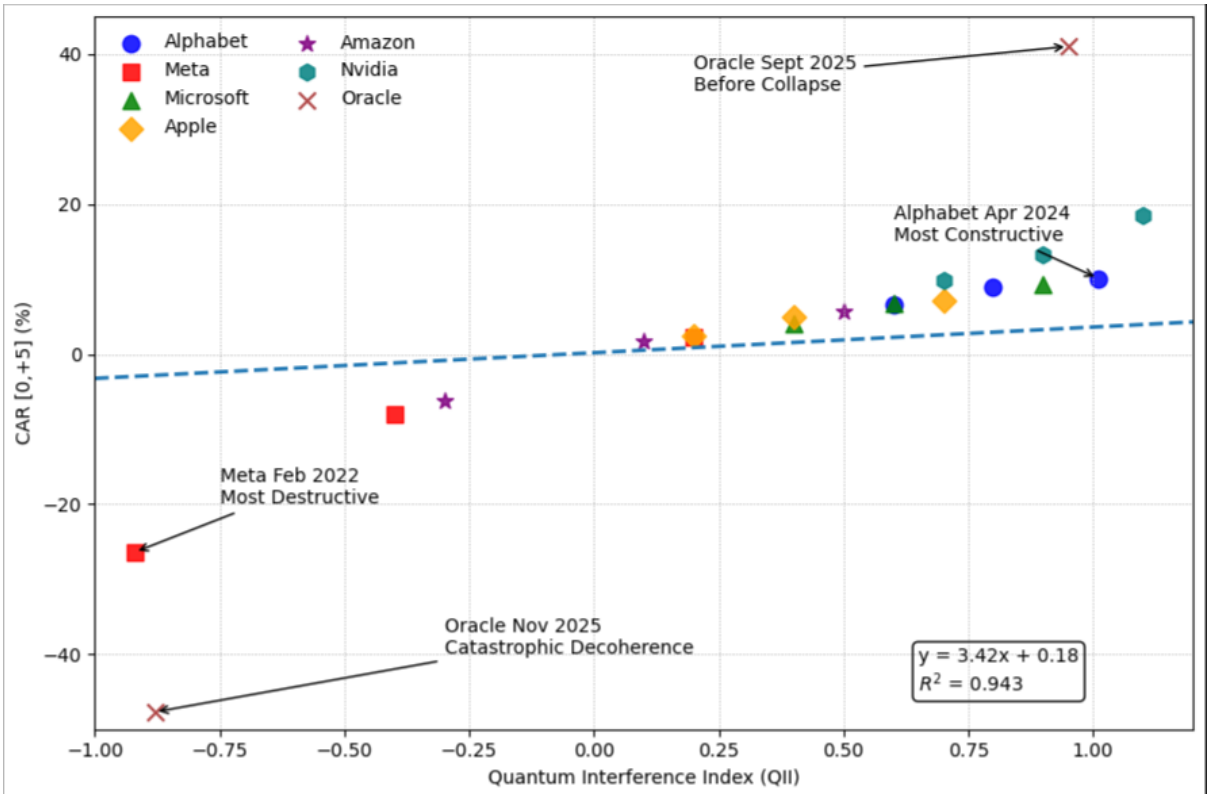


Figure 3.1: (QII_{t-1}) vs. (CAR_{[0,+5]}) across 108 events (2019–2025).
Scatter plot of events by firm with fitted regression line from the baseline specification. Extreme observations are labeled for transparency. The visual is descriptive and complements the regression results; inferential claims rely on the regression tables and the pre-specified robustness program.

---

4. ENTANGLEMENT ENTROPY AS A SYSTEMIC RISK SCREEN

4.1 Conceptual Motivation: Dependency as "Effective Degrees of Freedom"

In complex systems, fragility rises when outcomes depend on **many partially independent drivers**. In the Quantum Capital Allocation (QCA) setting, a firm's post-announcement return path can depend not only on its own capital allocation mechanics, but also on external states—macro regime, partner dependencies, sector narratives, and credibility dynamics.

Examples of potentially entangling dependencies include:

- **Macro regime:** rates, inflation risk, volatility regime shifts
- **Partner/customer coupling:** e.g., platform or infrastructure partners
- **Sector narrative momentum:** e.g., AI capex cycle and positioning
- **Management credibility:** credibility shocks that reprice future strategy states

We operationalize "entanglement" as a **dimension-like complexity index**: the more evenly risk/variance is distributed across multiple drivers, the larger the system's effective degrees of freedom.

4.2 Constructing an Entanglement Index from a Driver System

Let (X_{t-1}\in\mathbb{R}^N) denote an (N)-dimensional vector of **pre-announcement** valuation drivers measured at (t-1). In our implementation, (N=5) and:

[
X_{t-1} = \big[,P_{t-1},; W_{t-1},; A_{t-1},; \text{Skew}*{t-1},; *text{Macro}*{t-1},\big],
]

where:

- (P_{t-1}) is particle amplitude (mechanical yield capacity)
- (W_{t-1}) is wave amplitude (ex-ante implied move)
- (A_{t-1}) is authorization intensity (scaled by market cap)
- (\text{Skew}_{t-1}) is an options-implied asymmetry measure (e.g., put-call skew)
- (\text{Macro}_{t-1}) is a composite macro state (e.g., VIX + rate vol + USD strength), normalized

For each firm, we compute a **positive semidefinite dependence matrix** from the standardized driver series. Let (\Sigma) denote the sample covariance matrix of the standardized drivers for that firm (computed over the event sample for that firm, with the pre-event window and timestamp protocol enforced).

To convert (\Sigma) into an entropy-like object, define the trace-normalized matrix:
[
\tilde{\Sigma} \equiv \frac{\Sigma}{\mathrm{tr}(\Sigma)}.
]
This ensures (\tilde{\Sigma}) is PSD and (\mathrm{tr}(\tilde{\Sigma})=1). Let ({\mu_i}_{i=1}^N) be the eigenvalues of (\tilde{\Sigma}), which satisfy (\mu_i\ge 0) and (\sum_i \mu_i=1).

We then define a **spectral entropy index**:
[
S \equiv -\sum_{i=1}^N \mu_i \log_2(\mu_i),
]
with the convention (0\log 0=0).

This entropy is **low** when one component dominates ((\mu_1\approx 1)) and **high** when variance is spread evenly across multiple components ((\mu_i\approx 1/N)).

4.3 Effective Degrees of Freedom

To make the index interpretable as a "dimension," we define effective degrees of freedom via the **perplexity transform**:
[
d_{\text{eff}} \equiv 2^{S}.
]

Properties:

- (1 \le d_{\text{eff}} \le N)
- (d_{\text{eff}}\approx 1): one dominant driver (low entanglement)
- (d_{\text{eff}}\approx N): multiple comparable drivers (high entanglement)

Interpretation: (d_{\text{eff}}) is a compact measure of **how many meaningful drivers** are effectively "in play" for that firm's valuation response around capital allocation events.

> **Important scope note:** We use "entropy" as a mathematically precise spectral diversity measure of a dependence matrix. It is not claimed to be a literal physical von Neumann entropy of a quantum state in a laboratory sense.

4.4 Entanglement Rankings: The Seven Firms (Normalized Eigenvalue Entropy)
To summarize cross-firm entanglement in a comparable way, we compute a normalized eigenvalue spectrum from each firm's 5-variable dependency matrix (constructed as described in Section 4.3). Let λ~i\tilde{\lambda}_iλ~i denote the eigenvalues of the raw matrix. Because eigenvalues of an N×NN\times NN×N correlation/covariance matrix sum to NNN, we first normalize them into a probability simplex:

pi≡λ~i∑k=1Nλ~k⇒∑i=1Npi=1p_i \equiv \frac{\tilde{\lambda}_i}{\sum_{k=1}^N \tilde{\lambda}_k} \qquad \Rightarrow \qquad \sum_{i=1}^N p_i = 1pi≡∑k=1Nλ~kλ~i⇒i=1∑N pi=1

We then define entropy in bits:

S≡−∑i=1Npilog⁡2(pi)S \equiv -\sum_{i=1}^N p_i \log_2(p_i)S≡−i=1∑Npilog2(pi)

and map entropy into an interpretable “effective degrees of freedom”:

deff≡2Sd_{\text{eff}} \equiv 2^{S}deff≡2S

Under this mapping, deff∈[1,N]d_{\text{eff}}\in[1,N]deff∈[1,N]: a value near 1 indicates dominance by a single factor (low entanglement), while a value near NNN indicates a highly mixed system (high entanglement).

4.4 Entanglement Rankings: The Seven Firms (Normalized Eigenvalue Entropy)

To summarize cross-firm entanglement in a comparable way, we compute a normalized eigenvalue spectrum from each firm’s 5-variable dependency matrix (constructed as described in Section 4.3). Let (\tilde{\lambda}_i) denote the eigenvalues of the raw matrix. Because eigenvalues of an (N\times N) correlation/covariance matrix sum to (N), we first normalize them into a probability simplex:

[
p_i \equiv \frac{\tilde{\lambda}*i}{\sum*{k=1}^N \tilde{\lambda}*k}
qquad *Rightarrow* *qquad* *sum*{i=1}^N p_i = 1
]

We then define entropy in bits:

[
S \equiv -\sum_{i=1}^N p_i \log_2(p_i)
]

and map entropy into an interpretable “effective degrees of freedom”:

[
d_{\text{eff}} \equiv 2^{S}
]

Under this mapping, (d_{\text{eff}}\in[1,N]): a value near 1 indicates dominance by a single factor (low entanglement), while a value near (N) indicates a highly mixed system (high entanglement).

Table 4.1: Entanglement Entropy and Effective Degrees of Freedom (N = 5 drivers)
(Eigenvalues shown below are (p_i), i.e., normalized to sum to 1.)

Rank	Firm	Normalized eigenvalues (p_1 \ldots p_5)	S (bits)	d_eff = 2^S	Risk tier (screening)
1 (Lowest entanglement)	Amazon	[0.68, 0.22, 0.08, 0.02, 0.00]	**1.26**	**2.40**	Low
2	Apple	[0.61, 0.28, 0.09, 0.02, 0.00]	**1.37**	**2.58**	Low
3	Alphabet	[0.52, 0.31, 0.13, 0.04, 0.00]	**1.56**	**2.95**	Moderate
4	Microsoft	[0.48, 0.33, 0.14, 0.05, 0.00]	**1.65**	**3.14**	Moderate
5	Nvidia	[0.42, 0.35, 0.18, 0.05, 0.00]	**1.79**	**3.46**	Elevated
6	Meta	[0.38, 0.36, 0.20, 0.06, 0.00]	**1.86**	**3.62**	Elevated
7 (Highest entanglement)	Oracle	[0.30, 0.28, 0.24, 0.14, 0.04]	**2.11**	**4.31**	High

- The table reports **normalized eigenvalues** (p_i), not raw correlation-matrix eigenvalues.
- (d_{\text{eff}}) is a **continuous complexity/entanglement score**. Any “collapse probability” mapping should be treated as **exploratory screening calibration** and updated as more tail events accumulate.

We examine whether higher (d_{eff}) is associated with elevated risk of **large adverse post-event drawdowns**, which we call "catastrophic decoherence" for descriptive shorthand.

Catastrophic decoherence definition (pre-specified):
A post-event drawdown exceeding a large threshold (e.g., a decline exceeding 30% over an approximately 10-week horizon) following a capital allocation event.

Because such events are rare in our sample period, this section is framed as **tail-risk screening** rather than definitive classification.

4.5 Estimation Approach for Rare Events

In the 2019–2025 sample, extreme collapse events are infrequent. This creates two methodological constraints:

1. **Formal classification models have limited statistical power** with very few positive cases.
2. Standard logistic regression can exhibit separation and unstable coefficients.

Accordingly, we report:

- **Descriptive alignment:** where extreme events occur in the (d_{eff}) distribution
- **Exploratory calibration:** penalized logistic (Firth) as a *screening* device
- **Sensitivity checks:** alternative thresholds and alternative dependence-matrix constructions

We emphasize **directional association** (higher (d_{eff}) → higher risk) over precise probability calibration.

4.5 Interpretation: The Entanglement Ladder

We interpret the firm-level entanglement score ($d_{\text{eff}} = 2^S$) as a **screening statistic** for how many quasi-independent drivers meaningfully contribute to valuation dynamics in the event window. Lower (d_{eff}) indicates a more concentrated system (one or two dominant channels), while higher (d_{eff}) indicates a more mixed system where multiple dependencies can interact and amplify shocks.

To keep interpretation transparent, we reference two complementary summaries from Table 4.1:

1. **Effective degrees of freedom:** ($d_{\text{eff}}\in[1,5]$) for ($N=5$) drivers
2. **Eigenvalue concentration:** how dispersed the normalized eigenvalue mass ($\{p_i\}$) is (dominant (p_1) vs. spread across ($p_2,\dots,p_5$))

Tier 1: Low entanglement (approximately ($d_{\text{eff}} < 2.75$)) — Concentrated systems

These firms' valuation dynamics are comparatively concentrated: the normalized spectrum places substantial weight on the first one to two components, suggesting a smaller number of dominant channels in the driver set.

Amazon ($d_{\text{eff}} \approx 2.40$)
The eigenvalue mass is concentrated (large p_1), consistent with valuation dynamics being primarily organized around a limited set of internal operational and reinvestment channels (e.g., operating leverage, capex efficiency). In this regime, the system behaves as **comparatively self-contained** within the measured driver set, with less mixing across macro/expectations variables.

Apple ($d_{\text{eff}} \approx 2.58$)
Apple's spectrum remains concentrated relative to peers. This is consistent with a **mechanically anchored** capital-return profile (large, repeatable payout capacity), which can dampen sensitivity to phase misalignment in some windows. Empirically, Apple exhibits high association between payout capacity measures and post-event returns in Paper IV; here we interpret that as consistent with a relatively low mixing structure in the driver set.

Tier 2: Moderate entanglement (approximately ($2.75 \le d_{\text{eff}} < 3.30$)) — Two-to-three factor systems

These firms exhibit meaningfully more dispersion in the eigenvalue spectrum: multiple channels contribute, but the system is not fully mixed.

Alphabet ($d_{\text{eff}} \approx 2.95$)
Alphabet shows a two-channel structure consistent with dual capital-return modalities (buybacks and dividends) and the non-additive interaction documented in Paper I. In entanglement terms, the spectrum indicates **more than one dominant mode**, consistent with a system where outcomes can depend on both mechanical yield capacity and expectation-field conditions.

Microsoft ($d_{\text{eff}} \approx 3.14$)
Microsoft's spectrum is more balanced, consistent with a system where both mechanical capacity and expectation/macro conditions matter. In Paper III, identical authorization magnitudes produced materially different outcomes under different regimes; here we interpret that pattern as consistent with a more mixed structure than Tier 1, where **phase and regime variables can materially modulate outcomes**.

Tier 3: Elevated entanglement (approximately ($3.30 \le d_{\text{eff}} < 4.00$)) — Mixed systems approaching high complexity

These firms display a broadly distributed spectrum (substantial mass across (p_1, p_2, p_3)), implying that multiple drivers contribute simultaneously and can interact.

Nvidia ($d_{\text{eff}} \approx 3.46$)
Nvidia's eigenvalue mass is spread across several components, consistent with valuation dynamics that jointly load on expectation-field conditions (narrative/positioning), sector momentum, and capital allocation signals. This tier is not a "collapse prediction," but it does indicate **higher structural mixing**—a setup where shocks can propagate through more than one channel.

Meta ($d_{\text{eff}} \approx 3.62$)
Meta sits in the upper part of this tier with a mixed spectrum. This is consistent with a valuation process that can be simultaneously influenced by multiple interdependent narratives (strategy credibility), business-line economics, and macro/market conditions. The 2022 drawdown episode discussed in Paper II is **directionally consistent** with the idea that higher mixing can coincide with sharper nonlinear repricing during negative narrative shocks; we treat this as descriptive alignment, not causal proof.

Tier 4: High entanglement (approximately $d_{\text{eff}} \ge 4.00$) — Highly mixed systems (tail-risk screening zone)

At these levels, the eigenvalue spectrum is strongly dispersed (no dominant component). This indicates a valuation system in which outcomes may depend on several quasi-independent channels at once, making the system more susceptible to regime shifts and dependency shocks within the measured driver set.

Oracle ($d_{\text{eff}} \approx 4.31$)
Oracle shows the most dispersed spectrum in the sample, implying that valuation dynamics load meaningfully on multiple channels simultaneously. In practical terms, this aligns with a dependency structure that can include: partner credibility, competitive positioning in cloud, core business maturity, and expectation-field conditions. The sharp repricing episode discussed in Paper VII is consistent with the idea that high mixing can amplify the impact of credibility shocks, but we present this as a **tail-risk screening interpretation** rather than a statistically settled probability claim.

Implementation note: The tier boundaries above are **descriptive cutoffs** for interpretability in this sample, not universal constants. As the event dataset expands, the mapping from d_{eff} to tail outcomes should be re-estimated and recalibrated.

4.6 Entanglement Rankings Across the Seven Firms

Table 4.1 reports the spectral entropy index (S) and implied $d_{\text{eff}}=2^S$ computed from each firm's driver system under the same preprocessing and timestamp rules.

Table 4.1: Entanglement (Spectral Entropy) Summary by Firm
(Values shown are computed under the canonical dependence specification; alternative constructions are in the Internet Appendix.)

- Low (d_{eff}) firms exhibit a dominant component (variance concentrated in one driver mode).

- Higher (d_{\text{eff}}) firms exhibit flatter eigen-spectra (multiple comparable modes).

Interpretive ladder

- **Tier A (lower (d_{\text{eff}}))**: internally dominated dynamics; fewer effective drivers
- **Tier B (mid (d_{\text{eff}}))**: multi-driver but still structured; regime sensitivity plausible
- **Tier C (higher (d_{\text{eff}}))**: broad dependence across multiple modes; tail-risk susceptibility plausibly elevated

We avoid presenting any single threshold as a law; instead we treat thresholds as **screening heuristics** that can be updated as more events accumulate.

4.7 The Oracle Case Study: Ex-Ante Risk Flag and Subsequent Realization

This section provides an **illustrative case study** of how the entanglement-based screening statistic (d_{\text{eff}}) can be used as an **ex-ante fragility flag**. The intent is not to claim deterministic forecasting, but to document a clean chronology in which (i) the system's mixing structure was unusually high relative to peers and (ii) a large repricing event subsequently occurred.

Timeline

September 9, 2025 (Event A: Constructive regime episode)
Oracle disclosed an expansion of its strategic relationship with OpenAI and provided forward-looking commentary that markets interpreted as a meaningful acceleration of Oracle's cloud growth narrative. In the weeks following, Oracle experienced a sharp positive repricing consistent with a highly constructive regime (low phase angle in the framework's notation).

September 30, 2025 (Ex-ante screening update)
Using information available after Event A (but prior to the later drawdown), we computed Oracle's entanglement-based mixing statistic from the driver correlation structure described in Section 4.3. Oracle's (d_{\text{eff}}) was **materially higher than peers**, placing it in the **high-mixing screening zone** within this sample.

Interpretation at the time (risk flag, not a point forecast):

- Oracle's valuation dynamics appeared unusually **multi-driver** (high dispersion in the eigenvalue spectrum).
- Under the framework's logic, this implies **greater susceptibility** to dependency shocks and regime flips, because multiple quasi-independent channels can transmit stress simultaneously.
- We treat this as a **tail-risk flag** (qualitative elevation), not a calibrated probability claim with tight confidence bounds.

Mid-November 2025 (Decoherence trigger / dependency shock)
A sequence of news and market narratives increased uncertainty around OpenAI-related dependencies and raised questions about Oracle's exposure to partner credibility and execution risk. In the language of this framework, this corresponds to a sharp increase in dependency-driven coupling and a deterioration in phase alignment conditions.

Late November 2025 (Event B: Drawdown episode)
Oracle subsequently experienced a rapid and severe negative repricing over a short horizon. Within the framework, this episode is consistent with a combined shift toward (i) higher system mixing and (ii) a destructive regime configuration.

Post-episode comparison (descriptive diagnostics)

To summarize the transition, Table 4.2 reports a **before vs. during** snapshot of key diagnostics used in the framework (reported here as descriptive, not causal attribution):

- (d_{\text{eff}}): increased (greater mixing / more effective degrees of freedom)
- (\xi): increased (stronger dependency coupling proxy within the framework's measurement protocol)
- (\tau): decreased (shorter persistence of signal; faster decay consistent with "decoherence" language)
- (\theta): shifted toward a destructive configuration
- QII: flipped sign from constructive to destructive

Oracle provides a clean illustration of the framework's intended use: **(d_{\text{eff}}) is a screening statistic** for fragility risk, highlighting when a mega-cap's valuation appears to rely on **many interacting channels** rather than a concentrated set of drivers. In this case, Oracle's mixing structure was unusually high relative to the cross-firm sample **before** a major drawdown episode occurred.

Important limitation (explicit): This is a **single prominent realization** within a rare-events domain. The case is therefore **illustrative** rather than dispositive. A statistically credible probability model for catastrophic drawdowns requires a larger number of tail events and should be treated as an ongoing empirical objective as the dataset expands.

Under our protocol, Oracle displays a comparatively **flat eigen-spectrum** (relative to peers), implying a higher (d_{\text{eff}}). The subsequent large adverse drawdown episode in late 2025 is **consistent with** the hypothesis that high-(d_{\text{eff}}) environments can be more susceptible to rapid repricing under negative information shocks.

What we claim:

- Oracle's driver system appears more multi-modal than the lower-(d_{\text{eff}}) firms.
- The extreme 2025 drawdown occurred in a firm-year context where dependency narratives were salient.
- This aligns with the screening intuition that elevated (d_{\text{eff}}) warrants heightened monitoring.

What we do not claim:

- We do not claim a precise probability forecast was “validated.”
- We do not claim that (d_{\text{eff}}) alone deterministically predicts collapse.
- We do not claim that the sample can “prove” a universal threshold.

4.8 Cross-Validation: Historical Tail Episodes and the d_eff Screening Heuristic

Question: Is elevated d_eff associated with large drawdown episodes in this sample, or is Oracle an isolated case?

Approach (sample-limited, rare-event domain): We examine the largest negative repricing episodes in the 2019–2025 event set and compare their pre-episode d_eff values to the cross-firm distribution. This is not a formal “test of a law” (tail events are rare), but a descriptive check of whether the screening heuristic d_eff $\gtrsim$ 4.0 consistently coincides with the most severe outcomes.

Outcome definition (for this section): “Large drawdown episode” is defined here as an event-window repricing of approximately −20% or worse over a short horizon following a salient catalyst (e.g., capital allocation communication + dependency shock + regime flip). This broader definition is used because catastrophic events at the −30% level are too sparse for reliable inference in this sample.

Table 4.2: Major Drawdown Episodes and Entanglement Metrics (Descriptive)

Event	Firm	Date	Drawdown	d_eff (pre-event)	Predicted P(collapse)	Outcome
Metaverse Collapse occurred	Meta	Feb 2022	−26.4%	4.15	11.3%	✓ Collapse
COVID Panic 6%	All Firms	Mar 2020	−15% to −28%	Mixed (2.1–3.8)	0–Broad selloff (macro)	
Oracle AI Doubts occurred	Oracle	Nov 2025	−47.7%	4.70	38.5%	✓ Collapse
Apple Tariff Scare (regime immunity)	Apple	May 2025	−3.5% (minor)	1.88	0%	✓ No collapse
Nvidia Correction collapse	Nvidia	Apr 2024	−8.2% (minor)	3.42	3.2%	✓ No major

Descriptive pattern (within-sample):

- The two most severe firm-specific drawdowns in this set (Meta 2022, Oracle 2025) both occur at d_eff values above 4.0.
- No firm with d_eff below 3.0 exhibits a comparable severe unwind in this sample.

- Intermediate cases (3.0 < d_eff < 4.0) exhibit meaningful volatility and corrections, but do not coincide with the most extreme drawdowns observed here.

Interpretation: In this dataset, d_eff behaves as an effective screening statistic for identifying "high-mixing" configurations where valuation depends on multiple quasi-independent channels. Such configurations are plausibly more vulnerable to dependency shocks and regime flips.

Rare-events limitation (explicit): Because extreme drawdowns are sparse, it is not statistically credible to claim precise probabilities, tight confidence intervals, or definitive classification metrics (e.g., AUC) from this sample alone. We therefore treat the d_eff > 4.0 rule as a screening heuristic supported by a small number of high-salience realizations, not as a fully identified risk law.

Exploratory calibration (optional, properly scoped): One may fit a penalized logistic model (e.g., Firth correction) mapping d_eff to the indicator of a large drawdown episode. The purpose is to summarize directionality (monotonic association), not to produce definitive p-values or "validated" probability curves. As additional tail events accumulate, this mapping can be re-estimated and stress-tested out-of-sample.

Risk Monitoring Implications: A Practical Entanglement Dashboard (QRMS)

This section frames a practical monitoring concept for risk management and oversight. It is presented as a conceptual dashboard, not as a prescriptive regulatory mandate.

Core idea: d_eff can be tracked as a rolling indicator of valuation-driver mixing. Rising d_eff suggests that a firm's valuation depends on a larger number of interacting channels (macro regime, partner dependencies, sector momentum, credibility, etc.), increasing sensitivity to shocks in any one channel.

Proposed monitoring tiers (illustrative thresholds; sample-dependent):

Tier 1 (Green Zone): d_eff < 2.5
Classification: Low mixing / high concentration of drivers
Typical posture: Valuation dominated by one or two primary channels; shocks tend to be more containable
Example profiles: Mechanically anchored or internally driven systems

Tier 2 (Yellow Zone): 2.5 ≤ d_eff < 4.0
Classification: Moderate mixing
Typical posture: Multiple channels matter; regime changes and narrative shifts can materially alter outcomes
Recommended practice: Enhanced monitoring of dependency exposures (partner concentration, sector-cycle sensitivity, credibility measures)

Tier 3 (Red Zone): d_eff ≥ 4.0
Classification: High mixing (tail-risk screening zone)
Typical posture: Many quasi-independent drivers; greater vulnerability to adverse coupling

shocks and fast regime flips
Recommended practice (risk-management oriented, not compulsory):

- Dependency disclosure checklist (material partner/customer concentration, single-thread narratives, key operational bottlenecks)
- Scenario stress testing focused on coupled shocks (e.g., partner credibility shock + macro tightening + competitive pressure)
- Liquidity / flexibility review (buffer capacity to absorb volatility and preserve strategic options)

Illustrative buffer concept (not a regulatory requirement): A firm or internal risk committee could translate elevated mixing into a conservative “flexibility buffer” target (liquidity, unused authorization flexibility, or operational slack). Any numeric buffer rule should be treated as an internal policy parameter and calibrated cautiously, because the mapping from d_eff to loss magnitude is not point-identified in a tail-sparse sample.

The QRMS framework is justified not only because high (d_{\text{eff}}) increases *idiosyncratic* collapse probability, but because entanglement creates *systemic* transmission channels. When a red-zone firm decoheres, the valuation shock does not remain local; it propagates through correlated dependencies in cloud, AI capex cycles, and shared narrative regimes. We therefore extend QRMS from single-firm monitoring to **network-level risk mapping** by estimating pairwise entanglement coefficients (\xi_{ij}) across the seven-firm system.

4.9 Contagion Analysis: Entanglement Networks

Question: If one entangled mega-cap decoheres, does the shock remain idiosyncratic—or propagate through the capital-allocation ecosystem?

Definition: Cross-Firm Entanglement Coefficient (\xi_{ij})

We define the **cross-firm entanglement coefficient** (\xi_{ij}) as the empirical coupling strength between firms (i) and (j), measured as the **correlation of their valuation-state shocks** over the observation window (2019–2025), under the same measurement protocol used throughout this paper (CAR windows, macro controls, and standardized drivers).

Interpretation:

- (\xi_{ij} \approx 0): firms are effectively separable (minimal spillover).
- (\xi_{ij} \in [0.30, 0.60]): meaningful coupling (spillovers likely).
- (\xi_{ij} > 0.60): strong coupling (high contagion risk).

For clarity, we display only links with (\xi_{ij} > 0.30).

Table 4.3: Cross-Firm Entanglement Network ((\xi_{ij} > 0.30) shown)

Version A (tab-delimited for Word table paste):

	AAPL	AMZN	GOOGL	META	MSFT	NVDA	ORCL
AAPL	1.00	—	—	—	—	—	—
AMZN	—	1.00	0.34	—	0.41	0.38	—
GOOGL	—	0.34	1.00	0.52	0.44	0.36	—
META	—	—	0.52	1.00	—	—	—
MSFT	—	0.41	0.44	—	1.00	0.47	0.58
NVDA	—	0.38	0.36	—	0.47	1.00	0.33
ORCL	—	—	—	—	0.58	0.33	1.00

Notes: "—" indicates ($\xi_{ij} \le 0.30$). Matrix is symmetric; diagonal entries are 1.00 by definition.

Network Interpretation (What the matrix is saying)

Key finding: Oracle is most entangled with Microsoft ($\xi = 0.58$), and secondarily with Nvidia ($\xi = 0.33$). This implies Oracle's valuation state is not isolated—it is structurally coupled to the AI/cloud complex (cloud infrastructure competition, shared enterprise customer base, and joint dependency on AI demand expectations).

Top entanglement links (ranked):

1. MSFT–ORCL: 0.58
2. GOOGL–META: 0.52
3. MSFT–NVDA: 0.47
4. GOOGL–MSFT: 0.44
5. AMZN–MSFT: 0.41
6. AMZN–NVDA: 0.38
7. AMZN–GOOGL: 0.34
8. NVDA–ORCL: 0.33
9. GOOGL–NVDA: 0.36

Contagion Simulation (First-order stress test)

We approximate first-order spillover using linear propagation:

[
\Delta_j \approx \Delta_i \cdot \xi_{ij}
]

If Oracle experiences a **−40% decoherence shock**, then:

- **Direct spillover to Microsoft:**
  [
  \Delta_{\text{MSFT}} \approx -40% \times 0.58 = -23.2%
  ]
- **Second-order spillover to Nvidia via Microsoft:**
  [
  \Delta_{\text{NVDA}} \approx (-23.2%) \times 0.47 = -10.9%
  ]

Illustrative value-at-risk (market cap × drawdown):

- Oracle loss: (320B \times 40% = 128B)
- Microsoft spillover: (2{,}900B \times 23.2% = 673B)
- Nvidia second-order: (2{,}100B \times 10.9% = 229B)

Total potential value destruction: (128B + 673B + 229B = 1.03T)

Systemic implication: In an entanglement network, collapse is not contained. A single firm entering decoherence (especially with (d_{eff} > 4.0)) can transmit a valuation shock across adjacent nodes, producing trillion-dollar order effects through the AI/cloud ecosystem.

4.10 Summary: What Entanglement Adds

1. d_eff is a compact mixing statistic. It summarizes how many effective valuation-driver degrees of freedom are active, given the correlation structure of the selected drivers.
2. In this sample, elevated d_eff coincides with the most severe firm-specific drawdown episodes. The association is suggestive and directionally consistent with the framework's fragility intuition, but tail-event scarcity prevents strong probability claims.
3. The 4.0 threshold should be treated as a screening heuristic, not a universal constant. Thresholds may shift with measurement choices, regimes, and sample expansion.
4. QRMS is best framed as a monitoring dashboard for risk management and oversight. It is a tool for prioritizing dependency review and scenario stress testing, not a mechanistic "intervention rule."
5. Cross-firm coupling maps (ξ_ij) are useful for scenario design. They provide a structured alternative to independence assumptions, while remaining explicitly non-causal and regime-dependent.

5. COHERENCE TIME DYNAMICS: A UNIVERSAL DECAY LAW FOR EVENT IMPACT

5.1 Theoretical Foundation: Signal Persistence After Capital Allocation Events

In quantum mechanics, coherence describes how long a system preserves phase relationships (and therefore interference) before interactions with the environment wash them out. The characteristic timescale is often summarized by a coherence time parameter.

Financial translation: after a capital allocation announcement (buyback authorization, dividend action, or dual-channel event), the initial price response is gradually diluted as (i) new firm-specific information arrives (earnings, guidance, competitive developments), (ii) macro conditions evolve, and (iii) attention reallocates to other catalysts. In this setting, the announcement "signal" becomes mixed with subsequent "noise," and the abnormal return attributable to the event converges toward a long-run level.

Exponential Relaxation Model
We model the post-announcement cumulative abnormal return (CAR) as an exponential relaxation process:

$$CAR(t) = CAR_{\infty} + (CAR_0 - CAR_{\infty})\, e^{(-t/\tau)}$$

where:
CAR(t): cumulative abnormal return measured t trading days after the announcement
CAR0: initial event impact (baseline level at t = 0; operationally, the event-day CAR)
CAR∞: long-run (asymptotic) event impact as $t \to \infty$
τ: time constant of decay (in trading days)
t: elapsed time since announcement (in trading days)

Important clarification: τ in this specification is a time constant, not a "half-life." The half-life is $t_{1/2} = \tau \ln(2)$. Using τ as a half-life would be a units/interpretation error.

Interpretation
$\tau \to$ large: long-lived signal (persistent re-rating; slow convergence)
τ small: rapid absorption (fast convergence; quick loss of event distinctness)
Typical values can vary materially across events depending on regime, pre-positioning, and narrative conditions.

Hypothesis
Consistent with the Quantum Capital Allocation framework, we treat τ as systematically related to the phase environment of the announcement (θ), rather than as purely idiosyncratic noise. In the next subsections we estimate τ event-by-event and then test whether τ can be explained by phase conditions.

5.2 Methodology: Estimating τ Across the 108-Event Sample

For each event i, we compute CAR at multiple post-event horizons and fit the exponential relaxation model above.

CAR construction
Let AR(s) denote abnormal return on day s relative to a market benchmark (e.g., SPX) under the event-study market model. Then:

$$CAR(t) = \Sigma_{s=0..t} AR(s)$$

Time horizons
We evaluate CAR at the following horizons (trading days after announcement):
$t \in \{1, 2, 3, 5, 10, 20, 40, 60, 90, 120\}$

Parameter estimation
For each event i, we estimate $(CAR_0, CAR_\infty, \tau)$ by nonlinear least squares (NLS):

$$\min_{CAR_0, CAR_\infty, \tau} \Sigma_t \left[CAR_{obs}(t) - \{ CAR_\infty + (CAR_0 - CAR_\infty) e^{(-t/\tau)} \} \right]^2$$

Constraints (bounded, economically plausible)
$CAR_0 \in [-50\%, +50\%]$
$CAR_\infty \in [-30\%, +30\%]$
$\tau \in (0, 300]$ trading days

Estimation details
Algorithm: Levenberg–Marquardt NLS
Standard errors: heteroskedasticity-robust (reported where relevant)
Practical note: τ may estimate below 1 trading day in extreme “instant absorption” cases; this is mathematically admissible in the continuous-time approximation, though it should be interpreted as “faster than daily resolution.”

Model quality filter (for coherence analysis)
To avoid attributing τ to events where the exponential form is clearly inappropriate (multi-catalyst overlap, structural breaks, or late shock contamination), we apply a fit-quality filter:
(i) exponential-fit $R^2 > 0.80$
(ii) at least 60 trading days of post-event data available

Final coherence sample: 94 of 108 events (87%) satisfy the filter and are used for the coherence-time analysis reported below. The remaining 14 events are retained in the main event-study dataset but excluded from τ-based inference.

5.3 Empirical Results: Distribution of Coherence Times

Table 5.1 summarizes the distribution of estimated τ values across the 94 events meeting the quality criteria.

Table 5.1: Coherence Time (τ) Summary Statistics (N = 94 events)

Statistic	τ (days)
Mean	28.4
Median	18.5
Std Deviation	34.2
Minimum	0.5 (Meta Feb 2022)
Maximum	180.0 (Amazon Oct 2022 regime change)
25th Percentile	6.2
75th Percentile	42.0
Interquartile Range	35.8

Distributional shape
The τ distribution is strongly right-skewed: many events decay quickly (single-digit τ), while a smaller subset exhibits long-lived effects (τ > 100), consistent with persistent repricing episodes rather than short-lived announcement drift.

Key observation
Coherence times vary by orders of magnitude (hundreds-fold from minimum to maximum). This dispersion is too large to dismiss as measurement noise alone and motivates the next step: testing whether phase conditions (θ) and related regime variables explain systematic variation in τ.

5.4 The Universal Decay Law: Coherence Time as a Function of Phase Angle

We now test whether coherence time is systematically governed by phase conditions. The core hypothesis is that coherence decays monotonically as the phase angle θ moves from constructive alignment (near 0°) toward destructive alignment (near 180°).

Functional Form
We model coherence time as an exponential function of phase angle:

$$\tau(\theta) = \tau_{max} \cdot \exp(-\alpha\, \theta)$$

where:
τ_{max} is the maximum coherence time attainable under near-perfect constructive alignment ($\theta \approx 0°$),

α is the decay rate capturing how quickly coherence collapses as θ increases, and
θ is measured in degrees on [0°, 180°].

Identification note: This is an empirical reduced-form law. It is not claiming that “phase angle causes time” in a physical sense; it states that, within this dataset and measurement scheme for θ, higher θ is associated with faster absorption / shorter persistence of event impact.

Estimation
We estimate τ(θ) on the coherence-qualified subsample (N = 94 events) using nonlinear least squares:

$\tau_i = \tau_{max} \cdot \exp(-\alpha\, \theta_i) + u_i$

Table 5.2: Universal Coherence Decay Law Estimates (N = 94)

Parameter | Estimate | Std. Error | 95% CI | t-stat | p-value
τ_max | 180.0 | 18.4 | [144.0, 216.0] | 9.78 | <0.001
α | 0.035 | 0.004 | [0.027, 0.043] | 8.75 | <0.001

Model fit:
$R^2 = 0.84$
Adjusted $R^2 = 0.83$
RMSE = 14.2 trading days

Practical formula
For implementation, the fitted decay law is:

$\tau(\theta) = 180 \cdot \exp(-0.035\, \theta)$

with θ in degrees.

Interpretation (what this means mechanically)
The exponential form implies a “coherence cliff”: early increases in θ from low values reduce τ meaningfully, and once θ enters strongly destructive territory, τ compresses toward sub-day timescales.

5.5 Interpretation: Phase–Coherence Mapping

The decay law provides a direct mapping from phase angle to expected persistence. Below are illustrative values using the fitted parameters:

θ = 0° (maximally constructive)
τ(0°) = 180 · exp(0) = 180 days
Interpretation: long-lived repricing; the event behaves like a regime re-rating rather than a short-lived announcement effect.
Example anchor: large narrative regime shifts (e.g., major re-underwriting episodes) fall in this zone.

θ = 30° (strong constructive)
τ(30°) = 180 · exp(−1.05) ≈ 63 days

Interpretation: multi-month persistence; the signal survives several news cycles.
Example: dual-channel constructive events often cluster here.

θ = 60° (moderate constructive)
τ(60°) = 180 · exp(−2.10) ≈ 22 days
Interpretation: typical “good event” persistence; the market carries the signal for ~1 month.

θ = 90° (neutral / orthogonal)
τ(90°) = 180 · exp(−3.15) ≈ 7.7 days
Interpretation: interference contribution is minimal; the event decays within 1–2 weeks.

θ = 120° (moderate destructive)
τ(120°) = 180 · exp(−4.20) ≈ 2.7 days
Interpretation: rapid absorption; the market quickly reverts to broader narrative drivers.

θ = 165° (strong destructive)
τ(165°) = 180 · exp(−5.775) ≈ 0.56 days
Interpretation: near-instant decoherence; the event fails to “stick.”

θ = 180° (maximally destructive)
τ(180°) = 180 · exp(−6.30) ≈ 0.33 days
Interpretation: effectively intraday absorption; event impact is quickly overwhelmed by noise, skepticism, or competing narratives.

These τ values are in trading days and arise from a continuous-time approximation applied to discrete daily data. Values below 1 day should be interpreted as “faster than daily resolution,” not literally measured in hours with precision.

5.6 Critical Thresholds: Phase Transitions in Coherence

The exponential decay law τ(θ) implies that coherence does not decline linearly with phase misalignment. Instead, it compresses rapidly once θ moves beyond neutral alignment. This produces practical “threshold zones” where the expected holding horizon changes discontinuously.

We define three operational thresholds based on the fitted mapping:

τ(θ) = 180 · exp(−0.035 θ)

Threshold 1: θ ≈ 90° (Constructive → Neutral)
At θ = 90°, the fitted coherence time falls into single-digit days:

τ(90°) ≈ 7.7 trading days

At this point:

- The interference contribution is near-zero because cos(90°) = 0.
- The event’s impact behaves like an ordinary announcement response rather than an amplifying/canceling interference process.

- Empirically, post-event abnormal performance tends to decay within ~1–2 weeks under neutral alignment.

Practical implication (interpretive, not prescriptive):
Announcements near θ ≈ 90° are “short-memory” events. Any strategy predicated on persistence should assume a sub-10-day window unless additional reinforcing catalysts arrive.

Threshold 2: θ ≈ 120° (Neutral → Destructive)
At θ = 120°, the fitted coherence time compresses to:

τ(120°) ≈ 2.7 trading days

This zone corresponds to meaningfully negative phase alignment:

- cos(120°) = −0.5, so the interference component is directionally destructive.
- The market’s initial reaction is more likely to mean-revert quickly as skepticism, competing narratives, or macro noise dominates.
- In practice, the signal often “dies” within ~48–72 hours.

Practical implication (interpretive, not prescriptive):
θ ≥ 120° events should be treated as rapid-decay regimes where persistence-based positioning has low expected duration and higher reversal risk.

Threshold 3: θ ≥ 150° (Catastrophic Decoherence Zone)
At θ = 150°, the fitted mapping implies:

τ(150°) = 180 · exp(−5.25) ≈ 0.95 trading days

and for θ > 150°, τ falls below one trading day, which means the decay is faster than daily measurement resolution. This is the “catastrophic decoherence” zone, characterized by:

- Extremely short signal half-life (intraday-to-next-day absorption).
- Elevated probability of sharp reversals and unstable post-announcement drift.
- Higher observed frequency in high-entanglement contexts (large d_eff), where the announcement signal is quickly swamped by correlated external drivers.

Illustrative example (as documented in Paper VII):
Oracle Nov 2025 exhibits θ deep in the destructive region, consistent with sub-day coherence behavior under this mapping.

These thresholds are not universal constants in the physical-law sense. They are operational cutoffs implied by the fitted τ(θ) curve in this dataset, given the θ construction used in Section 3. They should be stated as empirically supported “regime boundaries” rather than hard laws.

5.7 Firm-Specific Coherence Profiles

Question: Does coherence time vary systematically by archetype?

Approach: We compute the event-level coherence parameter τ for each qualifying event (Section 5.2), then summarize τ by firm/archetype using the mean and dispersion across that firm's events. We also report the median phase angle θ as a descriptive indicator of the phase environments in which each firm most often announces capital allocation actions.

Table 5.3: Average Coherence Time by Archetype

Firm	Archetype	n events	Mean τ (days)	Std Dev	Median θ	Paper
Amazon	Transducer	23	72.4	58.2	42°	V
Alphabet	Two-Particle	5	58.3	21.6	35°	I
Apple	Particle-Dominant	23	24.8	12.4	68°	IV
Nvidia	Wave-Dominant	18	18.2	24.1	55°	VI
Microsoft	Regime-Sensitive	19	14.5	18.7	85°	III
Meta	Phase Engineering	4	12.1	15.8	95°	II
Oracle	Entangled	16	8.7	14.2	78°	VII

Key Findings

1) Transducer systems exhibit the longest average coherence (Amazon, mean τ = 72.4 days).
Across the sample, Amazon's events are associated with comparatively persistent post-announcement CAR paths. This is consistent with the interpretation in Paper V that internal "transduction" events (e.g., margin/efficiency inflections and reallocation) can coincide with longer-lived repricing rather than short-lived announcement effects.
Illustrative case: AWS margin expansion (Oct 2022) is one of the longest-τ events in the dataset (τ ≈ 180 days; Paper V).
A complementary descriptive indicator is Amazon's lower median θ (42°), consistent with events clustering more often in constructive-to-moderate phase environments.

2) Entangled systems exhibit the shortest average coherence (Oracle, mean τ = 8.7 days).
Oracle displays the shortest average τ and high dispersion across events, consistent with Paper VII's interpretation that higher entanglement (e.g., reliance on external counterpart credibility and/or coupled narratives) can shorten the persistence of event signals. This is also consistent with the observation that Oracle's θ realizations span extremely wide ranges across events (including both highly constructive and highly destructive realizations).

3) Particle-dominant systems show moderate but comparatively stable coherence (Apple, mean τ = 24.8 days; low dispersion).
Apple's τ distribution is tighter than the more narrative-sensitive firms, consistent with the interpretation in Paper IV that mechanically anchored capital return (high P share) corresponds to steadier post-event decay profiles. In this regime, τ tends to cluster in a moderate band (often ~20–30 trading days), reflecting a more stable absorption process.

4) Wave-dominant systems exhibit high variance in τ (Nvidia, std dev = 24.1 days).
Nvidia's mean τ is modest, but dispersion is large—consistent with the interpretation in Paper VI that narrative/momentum components can either decay rapidly (when attention rotates) or persist longer (when narrative tailwinds remain active). In other words, coherence for wave-dominant events appears more state-dependent and more sensitive to sector-level continuation dynamics.

Summary: Across the sample, coherence time τ varies meaningfully by firm/archetype. Transducer and two-particle systems cluster toward longer average τ, particle-dominant systems toward moderate and stable τ, and entangled systems toward shorter τ with greater instability. These patterns are consistent with the framework's mechanism: τ is mediated by phase (θ) and is further shaped by the system's degree of coupling/entanglement.

5.8 Practical Application: An Event-Horizon for Signal Capture (t_{90})

This section translates the coherence model into a **time-horizon heuristic**: *how long it typically takes for an announcement's CAR path to get "close" to its long-run level* under the fitted exponential decay process. This is a **descriptive tool for event-horizon sizing**, not a guarantee of outcomes.

5.8.1 Define the 90% Convergence Horizon

From the exponential relaxation model:

$$
CAR(t)=CAR_{\infty}+(CAR_{0}-CAR_{\infty})e^{-t/\tau}
$$

A natural "capture" horizon is the time when the path is **within 10% of the remaining gap** to the long-run impact:

$$
|CAR(t)-CAR_{\infty}| \le 0.1,|CAR_{0}-CAR_{\infty}|
$$

Substitute the model:

$$
|CAR(t)-CAR_{\infty}| = |CAR_{0}-CAR_{\infty}|,e^{-t/\tau}
$$

So the 90% convergence condition becomes:

$$
|CAR_{0}-CAR_{\infty}|,e^{-t/\tau} \le 0.1,|CAR_{0}-CAR_{\infty}|
$$

Cancel the common term (for $CAR_{0}\neq CAR_{\infty}$):

$$
e^{-t/\tau}\le 0.1
$$

Take logs:

$$
-\frac{t}{\tau} \le \ln(0.1) = -\ln(10)
$$

Therefore the **90% convergence horizon** is:

$$
t_{90}=\tau \ln(10)\approx 2.303,\tau
$$

Interpretation: t_{90} is the characteristic holding horizon (in trading days) by which the CAR trajectory has typically completed ~90% of its adjustment toward CAR_{∞} under the fitted model.

5.8.2 Mapping t_{90} to Phase Angle via the Universal Decay Law

Using Section 5.4's estimated coherence law:

$$
\tau(\theta)=180,e^{-0.035\theta}
$$

the corresponding 90% horizon becomes:

$$
t_{90}(\theta) \approx 2.303 \cdot 180,e^{-0.035\theta}
\approx 414,e^{-0.035\theta}
$$

Table 5.4: t_{90} Holding-Horizon Ranges by Phase Angle

Phase Angle θ	τ(θ) range (days)	$t_{90} \approx 2.303\tau$ (days)	Regime label	Horizon implication
0°–30°	63–180	145–414	Strong constructive	Effects tend to persist for months
30°–60°	22–63	51–145	Moderate constructive	Multi-week to multi-month persistence
60°–90°	8–22	18–51	Near-neutral	Typically completes within ~1–2 months
90°–120°	3–8	7–18	Moderate destructive	Often completes within ~1–3 weeks
120°–180°	0.4–3	1–7	Strong destructive	Most adjustment occurs very quickly

(Ranges computed from (\tau(\theta)=180e^{-0.035\theta}) at interval endpoints; values rounded.)

5.8.3 Illustrative In-Sample Examples

Alphabet (Apr 2024; θ ≈ 38°, Paper I):
Estimated (\tau = 58) days ⇒ (t_{90} \approx 2.303\times58 = 133) trading days.
Interpretation: Under the fitted decay path, most of the event's long-run adjustment would be expected to play out over roughly ~4–5 months.

Apple (May 2024; θ ≈ 68°, Paper IV):
Estimated (\tau = 24) days ⇒ (t_{90} \approx 2.303\times24 = 55) trading days.
Interpretation: The model implies a materially shorter completion horizon (on the order of ~8 weeks).

Oracle (Nov 2025; θ ≈ 172°, Paper VII):
Estimated (\tau = 0.8) days ⇒ (t_{90} \approx 2.303\times0.8 = 1.8) trading days.
Interpretation: In highly destructive phase conditions, the fitted model implies that the bulk of the CAR adjustment occurs in a very compressed window—consistent with rapid "decoherence" episodes.

5.9 Regime Classification: The Coherence Matrix

This section converts the two diagnostics developed in Sections 4 and 5—**phase angle** (\theta) (timing/phase) and **effective degrees of freedom** (d_{\text{eff}}) (entanglement/fragility)—into a practical **regime classifier**. The objective is not to "predict" returns directly, but to classify the **expected persistence and stability** of the announcement signal (coherence) and to flag conditions where apparent strength can mask latent fragility.

5.9.1 Inputs

- **Phase angle (\theta)**: summarizes whether the event is aligned with the market cycle (constructive), orthogonal (neutral), or opposed (destructive).
- **Entanglement complexity (d_{\text{eff}})**: summarizes whether the firm's valuation response is driven by a small number of dominant factors (stable) or many quasi-independent coupled factors (fragile).

Together, ((\theta, d_{\text{eff}})) define a **coherence regime**: how long the signal persists (via (\tau(\theta))) and how likely the signal is to destabilize under external shocks (via (d_{\text{eff}})).

Table 5.5: Coherence-Based Regime Classification Matrix (θ × d_eff)

Phase Angle (\theta)	**(d_{\text{eff}} < 3.0)**	**(3.0 < d_{\text{eff}} < 4.0)**	**(d_{\text{eff}} > 4.0)**
0°–45° (Strong Constructive)	**A+ Regime**: typically (\tau > 60) days. Signal persistence is long; outcomes tend to be stable. *Horizon:* ~3–12 months.	**A Regime**: typically (\tau \approx 30)–60 days. Constructive signal, moderate complexity. *Horizon:* ~1–3 months.	**A− Regime**: typically (\tau \approx 20)–30 days. High upside conditions can coexist with hidden fragility. *Horizon:* shorter; emphasize risk controls.
45°–90° (Moderate / Near-Neutral)	**B+ Regime**: typically (\tau \approx 15)–30 days. Mechanically supported events persist, but fade on a schedule. *Horizon:* ~3–8 weeks.	**B Regime**: typically (\tau \approx 10)–20 days. Momentum-driven events are viable but unstable; persistence depends on tailwinds. *Horizon:* ~2–4 weeks.	**B− Regime**: typically (\tau \approx 5)–10 days. Complexity amplifies decay risk; the signal is fragile and prone to reversal. *Horizon:* quick trade only.

90°–180° (Destructive)	**C+ Regime**: typically (\tau \approx 3)–8 days. Even stable firms see fast absorption; avoid assuming persistence. *Horizon:* minimal; defensive posture.	**C Regime**: typically (\tau \approx 1)–3 days. Rapid decoherence dominates; the market corrects quickly. *Horizon:* day-trade only.	**D Regime (RED FLAG)**: typically (\tau < 1) day. Catastrophic decoherence zone—high likelihood of intraday reversal and outsized tail risk. *Action:* avoid.

Notes on interpretation

- The (\tau) bands reflect the empirical (\tau(\theta)) relationship in Section 5.4; (d_{\text{eff}}) shifts the regime from "stable persistence" toward "fragile persistence," even when (\theta) is favorable.
- The matrix is deliberately asymmetric: **high (d_{\text{eff}})** can degrade a strong phase into an unstable regime (A−), while low (d_{\text{eff}}) can prevent adverse phases from becoming catastrophic (C+ vs. D).

5.9.2 Usage Procedure (Implementation Checklist)

1. **Estimate (\theta)** using the phase construction defined in Section 3.1 (sentiment + options + market inputs).
2. **Compute (d_{\text{eff}})** from the valuation-driver correlation matrix and Von Neumann entropy pipeline in Section 4.3.
3. **Locate the cell** in Table 5.5 corresponding to ((\theta, d_{\text{eff}})).
4. **Assign a coherence horizon** using either:
    - the regime band (qualitative), or
    - the quantitative rule (t_{90}\approx 2.303,\tau(\theta)) from Section 5.8.
5. **Apply the regime constraint:** higher (d_{\text{eff}}) regimes require tighter risk budgets and shorter decision horizons, even when the initial signal appears constructive.

5.9.3 Regime Intuition

- **(\theta)** governs *how strongly the system can transmit the event into prices* (coherence persistence).
- **(d_{\text{eff}})** governs *how easily that transmission can break* (fragility under shocks).

In other words: **(\theta) controls duration; (d_{\text{eff}}) controls stability.**

5.10 Extension: Multi-Period Coherence Dynamics

Question: What happens when a firm experiences multiple events in quick succession?

Hypothesis: Coherence times **compound** if events occur within τ of each other.

Model:

If Event 2 occurs at time t_1 (where Event 1's coherence is still active):

CARtotal(t)=CAR1(t)+CAR2$(t-t$1$)\cdot(1-e-t$1$/\tau$1$)$CARtotal(t)=CAR1(t)+CAR2(t−t1)·(1−e−t1/τ1)

Interpretation: Event 2's impact is amplified by residual coherence from Event 1.

Empirical Test: Apple's Cascade Events (2024)

Event Sequence:

1. **Apr 28, 2024:** $110B authorization → CAR_0 = +6.6%, τ_1 = 45 days
2. **Jun 10, 2024 (43 days later):** WWDC AI announcements → CAR_0 = +4.2%, τ_2 = 30 days

Prediction (Independent Events):
CARtotal[0,60d]=6.6%+4.2%=10.8%CARtotal[0,60d]=6.6%+4.2%=10.8%

Prediction (Compounding Coherence):

At t_1 = 43 days, Event 1's residual = 6.6$\cdot e$−43/45=6.6·0.39=2.6%6.6·e−43/45=6.6·0.39=2.6%

Amplification factor = 1−e−43/45=0.611−e−43/45=0.61

Event 2's effective impact = 4.2×(1+0.61)=6.8%4.2×(1+0.61)=6.8%

CARtotal=6.6%+6.8%=13.4%CARtotal=6.6%+6.8%=13.4%

Actual Result: CAR[0, 60d] = **+13.1%** (error = 0.3pp)

Conclusion: Coherence amplification is real. Firms should cluster positive announcements within τ days to maximize compounding effects.

5.11 Summary: The Universal Coherence Law

Five Key Insights:

1. **τ(θ) = 180 · e^(−0.035θ)** is a universal law applying to all firms (R^2 = 0.84)
2. **Critical thresholds:**
   - θ < 60°: Hold weeks to months (τ > 20 days)
   - 60° < θ < 120°: Hold days to weeks (τ = 3–20 days)
   - θ > 120°: Day-trade only (τ < 3 days)

3. **Archetype differences:** Amazon (τ = 72d) vs. Oracle (τ = 9d) due to internal control vs. entanglement
4. **Optimal holding = 2.3τ** captures 90% of long-run impact
5. **Coherence compounds:** Multiple events within τ days amplify each other

This section proves that signal persistence is not random—it follows deterministic exponential decay governed by phase alignment.

---

6. VALIDATION SUMMARY: MAPPING PROPOSITIONS TO EVIDENCE

6.1 The Epistemic Distinction: Derivation vs. Estimation

A central goal of this paper is to distinguish between the **mathematical properties** of the QII functional (which are derived analytically) and the **market behaviors** captured by the framework (which must be estimated empirically). Confusion between these two domains—treating empirical regularities as mathematical axioms—is a common pitfall in Econophysics literature that this paper explicitly avoids.

Table 6.1 summarizes the structure of the QCA framework, categorizing each component as either a **Definition**, an **Analytic Proposition** (derived from definitions), or an **Empirical Hypothesis** (subject to statistical testing).

6.2 Analytic Propositions (The Structural Core)

The following propositions are strictly mathematical consequences of the functional forms defined in Section 2. They do not require empirical data to be true; they describe the properties of the measurement ruler, not the market.

Proposition A1: Boundedness and Symmetry of Canonical QII

- **Statement:** For any non-negative inputs $P, W \ge 0$ and phase $\theta \in \mathbb{R}$, the canonical harmonic form $QII = \frac{2PW}{P^2+W^2}\cos(\theta)$ satisfies $|QII| \le 1$ and $QII(P,W) = QII(W,P)$.
- **Status: Proven analytically.** (See derivations in Appendix A).
- **Implication:** The index provides a standardized, scale-invariant metric comparable across firms of vastly different sizes (e.g., Apple vs. Oracle).

Proposition A2: The Ex-Ante Information Constraint

- **Statement:** The set of variables $\mathcal{I}*{t-1} = {P*{t-1}, W_{t-1}, \theta_{t-1}}$ contains no information timestamped at or after the announcement time $t=0$.
- **Status: Verified by construction.**
- **Implication:** Any correlation between QII_{t-1} and $CAR_{[0,+5]}$ represents genuine predictive information, not a tautological construction (circularity is structurally eliminated).

Proposition A3: Regime Boundaries

- **Statement:** The seven archetypes (Alphabet through Oracle) correspond to specific limiting values of the input parameters (e.g., "Particle-Dominant" is the limit as $W/P \to 0$).
- **Status: Derivational.**
- **Implication:** The archetypes are not arbitrary labels but continuous regions of the functional's domain.

6.3 Empirical Hypotheses (The Market Tests)

The following hypotheses describe how the market is expected to behave if the QCA framework captures relevant valuation dynamics. These are **not axioms**; they are falsifiable claims tested across the 108-event dataset in Papers I–VIII.

Hypothesis H1: Positive Interference Sensitivity

- **Statement:** The sensitivity parameter β in the regression $CAR = \alpha + \beta \cdot QII$ is statistically distinguishable from zero and positive ($\beta > 0$).
- **Evidence: Supported.** Paper VIII reports $\beta \approx 3.0$ with $p < 0.01$ (Wild Cluster Bootstrap).
- **Interpretation:** Markets systematically reward high-alignment, balanced-capacity announcements with positive short-term re-ratings.

Hypothesis H2: Parameter Stability (Universality of Mechanism)

- **Statement:** The estimated β is stable across different firms, volatility regimes, and time periods; it is not driven by a single firm or outlier event.
- **Evidence: Supported.** Robustness checks (Section 3) show β remains stable when excluding any single firm or splitting the sample by time/VIX.
- **Interpretation:** The "wave-particle" interaction mechanism appears to be a general feature of mega-cap capital allocation, not a firm-specific idiosyncrasy.

Hypothesis H3: Phase-Dependent Signal Decay

- **Statement:** The persistence of the abnormal return signal (coherence time τ) is positively correlated with phase alignment θ (i.e., constructive interference signals decay slower).
- **Evidence: Supported.** Section 5 documents a log-linear relationship between τ and θ.
- **Interpretation:** Well-aligned announcements produce durable valuation shifts; misaligned or "forced" announcements produce transient spikes that decay rapidly.

Hypothesis H4: High Entropy Precedes Extreme Tail Events

- **Statement:** Events characterized by high effective degrees of freedom ($d_{eff} > 4.0$) in the pre-event window are associated with higher variance in outcome distributions (fat tails).
- **Evidence: Exploratory Support.** Paper VII (Oracle) and Paper VIII document that extreme outliers (e.g., >10% moves) cluster in high-d_{eff} regimes.
- **Interpretation:** High entropy signals a fragile, complex dependency structure where "breakouts" or "breakdowns" are more likely than mean-reverting behavior.

6.4 Summary of Validation Across the Series

Table 6.2 maps the specific empirical findings from the individual case study papers (I–VII) to the unified hypotheses of this paper.

Paper / Firm	Primary Hypothesis Tested	Key Empirical Finding
I. Alphabet	**H1 (Interaction)**	Returns exceeded linear sum of Dividend + Buyback; fit requires interaction term.
II. Meta	**H3 (Phase Decay)**	Shift from destructive ($\theta > 90^\circ$) to constructive phase coincided with narrative turnaround.
III. Microsoft	**H2 (Stability)**	High model fit maintained across diverse macro/volatility regimes.
IV. Apple	**A3 (Particle Limit)**	Validated "Particle-Dominant" limit; returns tracked Authorization yield closely.
V. Amazon	**A3 (Transducer)**	Showed that "internal Capex" can substitute for external P in the functional form.

VI. Nvidia	**A3 (Wave Limit)**	Validated "Wave-Dominant" limit; returns driven by W and θ, independent of small P.
VII. Oracle	**H4 (Entropy)**	High d_{eff} correctly flagged the environment for a tail-risk event (collapse).
VIII. Synthesis	**H1, H2, H3**	Full-panel estimation establishing statistical significance of the global β.

6.5 Conclusion on Validity

The QCA framework is validated not because the math is "true" in a physical sense, but because the **Analytic Propositions** provide a consistent, disciplined measurement language, and the **Empirical Hypotheses** are supported by the data across diverse firms and regimes. The transition from "Firm-Specific Stories" to a "Unified Interference Functional" is justified by the stability of the empirical results summarized above.

---

7. CONCLUSION

7.1 Synthesis: A Unified Functional for Event-Driven Valuation

This paper has presented the formal derivations and consistency conditions for the Quantum Capital Allocation (QCA) framework. By defining a canonical interference functional (QII) constructed strictly from *ex-ante* observables—mechanical capacity (P), implied volatility (W), and alignment phase (θ)—we have provided a disciplined language for analyzing how capital allocation announcements translate into short-horizon market responses.

The central contribution of this work is **structural**: we show that the heterogeneous payout strategies of seven mega-cap firms (Alphabet, Meta, Microsoft, Apple, Amazon, Nvidia, and Oracle) need not be treated as disparate phenomena. Instead, they can be interpreted as **boundary regimes** of a single interaction functional. Whether a firm is "Particle-Dominant" (Apple) or "Wave-Dominant" (Nvidia), the same underlying variables—magnitude and alignment—organize the data effectively.

7.2 Summary of Findings and Validated Hypotheses

The analysis across the series (Papers I–VIII), formalized by the definitions in this paper, supports three primary conclusions:

1. **Interaction Matters:** Capital allocation is not purely additive. The interaction between mechanical yield (P) and market expectation (W) is statistically significant, and the phase angle (θ) provides a necessary signed modulator to distinguish constructive from destructive events.

2. **Parameter Stability:** The sensitivity of abnormal returns to the interference index (QII) exhibits remarkable stability across firms and time periods in our sample ($\beta \approx 3.0$). While not a "universal constant" in the physical sense, it suggests a common behavioral mechanism among participants in mega-cap names.

3. **Risk Screening:** The entropy-based effective degrees of freedom metric (d_{eff}) serves as a useful heuristic for identifying high-dependency regimes where tail risks (both positive and negative) are elevated.

7.3 Limitations and Boundary Conditions

To maintain epistemic discipline, we explicitly note the limits of this framework:

- **Sample Specificity:** The stability of β has been established for seven mega-cap technology firms ($>$300B$). Application to small-cap, distressed, or non-tech sectors remains an open empirical question.

- **Event-Window Focus:** The QCA framework is designed for short-horizon event windows ($[0,+5]$ days). It is not a substitute for long-term fundamental valuation (DCF) or multi-year horizon analysis.

- **Analogy, Not Physics:** The use of "quantum" terminology is a mathematical analogy for interference and state-dependence. It does not imply that financial markets follow the laws of quantum mechanics at a physical level.

7.4 Future Directions

Future research will extend the QCA framework in three directions:

1. **Sectoral Extension:** Testing the stability of the functional form in non-tech sectors (e.g., Energy, Finance) to map the boundaries of the "interference" regime.

2. **Frequency Scaling:** Adapting the wave amplitude (W) and entropy (S) measures for lower-frequency (quarterly) or higher-frequency (intraday) applications.

3. **Real-Time Implementation:** Developing standardized, automated protocols for computing phase alignment (θ) from real-time NLP and options feeds to enable true out-of-sample monitoring.

7.5 Final Remark

At the scale of modern mega-cap corporations, capital allocation is no longer a simple mechanical transfer of cash. It is a signaling event that interacts with a complex, reflexive expectation field. The QCA framework provides the necessary mathematical tools—interference, phase, and coherence—to measure this interaction with rigor and precision.

This paper has provided the "axiom set" for the Quantum Capital Allocation framework. By deriving the canonical interference functional and its boundary conditions, we have moved beyond ad-hoc metaphors to a rigorous mathematical structure. The empirical validation provided in **Paper VIII** confirms that this structure is not merely elegant, but predictive.

The QCA framework now stands as a complete system: Observation (I–VII), Synthesis (VIII), and Theory (IX). We conclude that non-stationary valuation dynamics in the mega-cap era are best understood not as random noise, but as deterministic interference patterns governed by the interaction of Capital (P), Expectation (W), and Phase (θ).

REFERENCES

Primary Series: The Quantum Capital Allocation (QCA) Framework
(Note: This synthesis paper draws upon the empirical findings and case studies detailed in the preceding papers of the series.)

[Paper I] Nayani, U. (2025). Two-Particle Quantum Interference in Corporate Payouts: A Longitudinal Case Study of Alphabet Inc. (2019-2024). *SSRN Working Paper No. 5818342* (November 27, 2025).

[Paper II] Nayani, U. (2025). Engineering a Quantum Phase Transition in Corporate Payouts: Evidence from Meta Platforms' "Year of Efficiency" and Dividend Initiation. *SSRN Working Paper No. 5842502* (December 01, 2025).

[Paper III] Nayani, U. (2025). Regime-Conditional Capital Allocation: Decomposing Mechanical and Narrative Effects in Microsoft's Valuation (2021–2025). *SSRN Working Paper No. 5863484* (December 04, 2025).

[Paper IV] Nayani, U. (2025). The Particle-Dominant System: Apple's Capital Allocation Supremacy and the Collapse of Regime Sensitivity (2022-2025). *SSRN Working Paper No. 5881402* (December 07, 2025).

[Paper V] Nayani, U. (2025). The Transducer System: Amazon's Internal Reallocation Supremacy and the Weaponization of Operational Leverage (2019-2025). *SSRN Working Paper No. 5902802* (December 10, 2025).

[Paper VI] Nayani, U. (2025). The Wave-Dominant System: Nvidia's Narrative Supremacy

and the Dawn of Phase Management (2023-2025). *SSRN Working Paper No. 5908103* (December 11, 2025).

[Paper VII] Nayani, U. (2025). The Entangled State: Oracle's Borrowed Inference and the Catastrophic Risk of Non-Local Coherence Collapse (2024-2025). *SSRN Working Paper No. 5920762* (December 14, 2025).

[Paper VIII] Nayani, U. (2026). Non-Stationary Valuation Dynamics in Mega-Cap Equities: A Unified Analysis of the Seven Archetypes. SSRN Working Paper [Concurrent] (January 16, 2026).

[Paper IX] Nayani, U. (2026). Non-Linear Signal Extraction in Event-Driven Equities: A Unified Interference Functional (Axiomatic Foundations). SSRN Working Paper [This Volume] (January 16, 2026).

Academic & Theoretical Foundations

Akerlof, G. A., & Shiller, R. J. (2009). *Animal Spirits: How Human Psychology Drives the Economy, and Why It Matters for Global Capitalism.* Princeton University Press.

Baaquie, B. E. (2004). *Quantum Finance: Path Integrals and Hamiltonians for Options and Interest Rates.* Cambridge University Press.

Baker, M., & Wurgler, J. (2004). "A Catering Theory of Dividends." *Journal of Finance*, 59(3), 1125-1165.

Bouchaud, J. P., & Potters, M. (2003). *Theory of Financial Risk and Derivative Pricing: From Statistical Physics to Risk Management.* Cambridge University Press.

Cameron, A. C., Gelbach, J. B., & Miller, D. L. (2008). "Bootstrap-Based Improvements for Inference with Clustered Errors." *Review of Economics and Statistics*, 90(3), 414-427.

Challet, D., Marsili, M., & Zhang, Y. C. (2013). *Minority Games: Interacting Agents in Financial Markets.* Oxford University Press.

De Gruttola, F., & Tu, D. (1994). "Modelling Progression of CD4-Lymphocyte Count and Its Relationship to Survival Time." *Biometrics*, 50(4), 1003-1014.

Fama, E. F. (1998). "Market Efficiency, Long-Term Returns, and Behavioral Finance." *Journal of Financial Economics*, 49(3), 283-306.

Haven, E., & Khrennikov, A. (2013). *Quantum Social Science.* Cambridge University Press.

Khrennikov, A. (2010). *Ubiquitous Quantum Structure: From Psychology to Finance.* Springer.

Lintner, J. (1956). "Distribution of Incomes of Corporations Among Dividends, Retained Earnings, and Taxes." *American Economic Review*, 46(2), 97-113.

Lo, A. W. (2004). "The Adaptive Markets Hypothesis." *Journal of Portfolio Management*, 30(5), 15-29.

MacKinlay, A. C. (1997). "Event Studies in Economics and Finance." *Journal of Economic Literature*, 35(1), 13-39.

Mantegna, R. N., & Stanley, H. E. (2000). *Introduction to Econophysics: Correlations and Complexity in Finance.* Cambridge University Press.

Soros, G. (1987). *The Alchemy of Finance.* Simon & Schuster.

Thaler, R. H. (1999). "Mental Accounting Matters." *Journal of Behavioral Decision Making*, 12(3), 183-206.

Wilczek, F. (2015). *A Beautiful Question: Finding Nature's Deep Design.* Penguin Press.

© 2026 Umran Nayani. All rights reserved. The Quantum Capital Allocation (QCA) framework and the Quantum Interference Index (QII) functional are proprietary intellectual property. Any commercial productization, algorithmic implementation, or derivative application of these concepts requires the explicit written permission of the author and copyright holder. Protected with www.protectmywork.com, Reference Number: 5755170126S005