



# 7<sup>th</sup> Dimension MARKETS

## THE RIPPLE EFFECT

*Multi-Regime Cross-Asset Impulse Modeling and Expected Alpha*

**Version 3.2 — April 2026**

7th Dimension Markets

### **Version 3.0 — New in This Release**

Adaptive Z-Score Windows — per-asset-class baseline calibration  
Nightly Automation Pipeline — data, shocks, X posts, and Substack across all 4 Worlds  
Nightly Conviction Cache — pre-computed walk-forward leaderboards stored in Supabase  
Forecast Tab Confidence Bands — full statistical profile: fan chart, outcome distribution histogram  
Capital Efficiency Backtester v2 — multi-driver logic, transaction costs, Monte Carlo validation  
Updated ADBE→NVDA Case Study — 2021–2026 results with transaction costs applied  
Platform Architecture — Free/Pro split, Google OAuth, Stripe subscription, TOS gate

### **Version 3.1 — April 2026 Update**

Confidence Score — corrected sample penalty exponent; documented three-factor formula  
Volume Confirmation — clarified dual-window approach (10d signal vs 20d chart overlay)  
Nightly Pipeline — migrated from GitHub Actions to Railway; corrected schedule reference  
World 2 Universe — expanded to 154 altcoins across 13 categories  
Backtester Defaults — clarified max hold cap vs. average hold period in case study  
Pricing — aligned with live website (\$99/\$79)

### **Version 3.2 — April 2026 Update**

Section 12: Real-Time Hourly Alert Pipeline — intraday Z-scores, time-scaling, live website widget  
Section 13: Z-Decay Exit Strategy — signal-aware exits, three-layer exit hierarchy, per-world defaults  
Section 14: Shock Clustering & Named Macro Events — correlation-based clustering, 12 event templates  
Renumbered Sections 12–15 to accommodate new content (previously 9–12)

## 1. Abstract (Executive Summary)

---

**The Premise:** Traditional correlation matrices fail because they measure everyday market noise. This paper introduces a Shock Engine that measures market reactions only during statistically significant volatility events.

**The Universe:** A four-world analysis spanning decades of Macro (World 1), modern Crypto (World 2), US Equity Giants (World 3), and GICS Sector ETFs (World 4) — 700+ assets tracked daily.

**The Finding:** Quantifiable proof that identifying isolated Z-score shocks in a Driver Asset yields highly predictable, tradable Alpha in Reaction Assets over a specific forward-looking time horizon. A trailing-stop event-driven strategy on the Adobe-to-Nvidia impulse generated a +405.2% net return (after costs) with a Sharpe ratio of 0.78, maintaining only 37.0% market exposure over a 5-year backtest period.

## 2. Methodology: The Regime-Conditioned Impulse Engine

---

Standard correlation matrices are fundamentally flawed for actionable trading because they measure average, everyday market noise. The Earth Scanner engine abandons continuous correlation in favor of an event-driven, discrete shock-detection framework. This model isolates extreme volatility events (the Impulse) in a specified Driver Asset and measures the subsequent directional drift (the Ripple) in a Reaction Asset over a defined time horizon.

### 2.1 Baseline Calibration and Shock Detection

To identify an anomaly, the engine must first dynamically define normal market conditions. The model employs a rolling statistical window (defaulting to 90 days) to calculate the localized mean and standard deviation of daily returns for each asset. A localized Z-score is calculated for every trading day:

$$Z = (R - \mu) / \sigma$$

Where R is the daily return,  $\mu$  is the 90-day rolling mean, and  $\sigma$  is the 90-day rolling standard deviation. An Impulse Event is registered only when an asset's daily return breaches a strict volatility threshold ( $|Z| \geq 1.5$  by default). The US Dollar Index is tracked via the UUP ETF (launched 2007), providing approximately 4,800 trading days of dollar shock history.

### 2.2 Lead-Lag Event Study (Expected Alpha)

Once a historical Impulse Event is isolated, the engine initiates a forward-looking event study. For every identified shock in the Driver Asset, the model tracks the cumulative return of the Reaction Asset over a specific Signal Horizon (default T+5 days). The Expected Alpha presented in the Impulse Matrix is the arithmetic mean of these subsequent returns across the entire historical sample:

$$E[\alpha] = (\text{Sum of Reaction Asset Returns over Horizon}) / N$$

Where N is the total number of historical shocks. A significantly positive Expected Alpha indicates that a shock in the Driver historically acts as a leading indicator for a rally in the Reaction Asset.

### 2.3 Walk-Forward Stability and Confidence Scoring

A high Expected Alpha is irrelevant if it is the result of a single black-swan outlier. To ensure the tradability of the signal, the engine executes a Walk-Forward Stability test. The historical dataset is partitioned into chronological folds. The model evaluates the directional consistency (Sign Rate) and magnitude retention (Mag Retention) of the Ripple Effect across different market regimes.

When walk-forward stability data is unavailable, the fallback formula applies:

$$\text{Confidence} = 100 \times (0.65 \times \text{SignRate} + 0.35 \times \text{MagRetention}) \times \text{FoldPenalty} \times \text{SamplePenalty}$$

When walk-forward stability data is available (the primary path for cached leaderboards), the engine uses a three-factor formula:

$$\text{Confidence} = 100 \times (0.45 \times \text{WF\_Stability} + 0.35 \times \text{SizeFactor} + 0.20 \times \text{EdgeFactor})$$

Where WF\_Stability is the normalized walk-forward stability score (0–1), SizeFactor is the sample penalty  $(1 - e^{-(N/20.0)})$  which saturates gradually as the event count grows, and EdgeFactor is the normalized absolute edge magnitude capped at  $|\text{edge}|/2.0$ .

The fold penalty  $(1 - e^{-(k/2.0)})$  ensures that leaderboards computed from only one or two folds are discounted. The Confidence Score (0–100) acts as a strict guardrail: relationships that score highly prove that the cross-asset ripple is a persistent structural market anomaly rather than a temporary statistical coincidence.

| Score Range | Tier            | Interpretation  |
|-------------|-----------------|---|
| 65 - 100    | High Confidence | Persistent structural anomaly, tradeable across regimes |
| 40 - 65     | Moderate        | Valid signal — verify with volume confirmation          |
| 0 - 40      | Low / Unstable  | Likely regime-specific or statistical coincidence       |

## 2.4 Intraday Data and Volume Confirmation

Earth Scanner v3 incorporates intraday OHLCV data (1-hour bars) for the most recent 30 days across all four worlds, stored in a dedicated hot-storage table alongside the daily archive. This enables a third dimension of signal quality beyond price Z-score and walk-forward stability: volume confirmation.

When a Z-score shock is detected, the engine compares today's intraday volume against a short-term baseline to identify institutional participation. Two volume windows serve different analytical purposes:

- Signal Confirmation (Forecast Tab): A 10-day moving average provides high sensitivity for detecting whether today's volume is abnormally elevated relative to the most recent trading activity. This is the primary signal quality gate.
- Visual Analysis (Macro Tab): A 20-day moving average — the standard “trading month” convention — provides a smoother overlay on price charts for visual trend identification.

An above-average volume reading on a shock day suggests institutional participation — the primary driver of sustained price moves. The Forecast tab surfaces this confirmation signal alongside the Expected Alpha day-by-day breakdown.

### Signal Quality Hierarchy

High Z-score + High Confidence Score + Above-Average Volume = Highest Conviction Setup

## 2.5 Adaptive Z-Score Windows

Version 3 introduces per-asset adaptive baseline calibration, replacing the single fixed window used in earlier versions. Different asset classes operate on fundamentally different volatility cycles — a 90-day baseline that accurately defines ‘normal’ for Gold or the S&P 500 is far too long for a mid-cap altcoin.

| Asset Class                               | Adaptive Window | Rationale   |
|---|-----------------|---|
| World 1: Macro (FX, Commodities, Indices) | 90 days         | Slow-moving fundamental drivers   |
| World 2: Alt-Crypto                       | 15 days         | Narrative-driven, high-volatility — short baseline prevents false positives |
| World 3: S&P 500 Universe                 | 90 days         | Earnings cycles align with quarterly horizon                                |
| World 4: Sector ETFs                      | 90 days         | Sector rotation operates on macro timescales                                |

The engine also applies a minimum data guard: if fewer than 20 data points are available for a given symbol's window, the system falls back to the maximum available history, preventing newly listed assets from generating false positives.

## 3. Exhibit A: Macro Foundations and the Liquidity Baseline (World 1)

---

Before analyzing nascent asset classes or specific equities, a quantitative model must establish a baseline of global liquidity. In the fiat system, liquidity is fundamentally governed by the relative strength of the United States Dollar (USD).

Using the Earth Scanner engine, we processed over 4,800 trading days of UUP (US Dollar Index ETF, 2007–present) alongside 15+ years of macro asset data to map the cross-asset ripples generated by shocks in the primary commodity and currency markets.

### 3.1 The US Dollar as the Global Denominator

The US Dollar Index (tracked via UUP) serves as the ultimate Driver Asset in the global macro system. Because commodities like Gold and Crude Oil are globally priced in dollars, and because massive amounts of international corporate debt are dollar-denominated, a sudden spike in USD strength acts as a vacuum for global liquidity.

### 3.2 The Drain Anomaly

By isolating extreme volatility events (Z-scores of 1.5 or greater) in the US Dollar Index over the full historical sample, the Impulse Matrix reveals a severe and persistent negative correlation across the broader macro universe. The engine identifies a strong negative Expected Alpha on standard risk assets and commodities in the 5-day Signal Horizon immediately following a USD upside shock.

This metric proves that the USD does not simply trend inversely to other assets on a daily basis; a shock in the dollar aggressively accelerates capital flight out of risk assets over the subsequent week.

### 3.3 Mechanism of Action

- **Commodity Deflation:** As the dollar strengthens, automated sell-offs in commodity futures are triggered
- **Risk-Off Cascades:** A sudden strong dollar signals tightening financial conditions, prompting institutional capital to liquidate equities and retreat into cash or short-term Treasuries
- **Cross-Currency Pressure:** Emerging market currencies weaken, amplifying the drain effect across global risk assets

## 4. Exhibit B: The Crypto Echo and Altcoin Beta (World 2)

---

While the US Dollar and commodities represent mature, highly liquid markets, the cryptocurrency ecosystem (World 2) exhibits a vastly different structural behavior. Nascent asset classes are heavily driven by speculative liquidity and narrative, resulting in extreme intra-market correlation — the Crypto Echo.

World 2 tracks 154 altcoins across 13 categories with full OHLCV history from 2020 to present, providing a statistically robust sample for impulse analysis across multiple crypto market cycles including the 2021 bull run, 2022 bear market, and 2023–2025 recovery. Categories include Layer 1/Layer 2 protocols, DeFi, GameFi, Meme, Interoperability, DePIN, Liquid Staking, and Real-World Assets (RWA).

### 4.1 The Global Impulse in Nascent Markets

In traditional equities, a shock in the Energy sector might have zero effect on Healthcare. In the cryptocurrency market, liquidity flows in a sequential waterfall: from fiat currency, into the largest market-cap coins (Bitcoin/Ethereum), and subsequently out into lower-cap altcoins. A Z-score shock in a primary driver asset rarely isolates itself — it triggers a Global Impulse where almost all altcoins move in lockstep.

### 4.2 The Altcoin Ripple and Amplification Detection

When isolating positive volatility shocks in Solana (SOL-USD), the Cross-Asset Impulse Matrix identifies a near-uniform Expected Alpha across the mid-cap altcoin space within a 5-day Signal Horizon. This tightly clustered return profile proves that mid-cap altcoins operate primarily as high-beta proxies for the major network tokens.

The true value of the matrix lies in identifying deviations from this baseline:

- Amplification (High Beta): Altcoins that historically return significantly more than the baseline — superior vehicles for capturing upside momentum
- Decoupling (Relative Weakness): Altcoins that return less or drift negative, indicating structural exhaustion or narrative loss

### 4.3 Tradable Takeaway

By utilizing the Confidence Score (Walk-Forward Stability) alongside the Expected Alpha, traders can systematically rotate capital into the altcoins that most reliably amplify the major market impulses, while avoiding those that are decoupling from the broader liquidity trend.

## 5. Exhibit C: Equity Giants and Sector Shocks (World 3)

While the US Dollar dictates global liquidity and Altcoins exhibit herd behavior, the US Equity market operates on a system of capitalization-weighted gravity. World 3 tracks the full S&P 500 universe (503 symbols) with daily OHLCV data from 2021 to present.

### 5.1 The Rising Tide Effect

At extreme volatility levels, mega-cap stocks act as macroeconomic drivers. When isolating positive volatility shocks in Apple, the historical data shows a near-uniform positive Expected Alpha across a diverse basket of S&P 500 peers over the subsequent 5-day Signal Horizon. This demonstrates the rising tide effect: a massive influx of capital into a market leader forces algorithmic rebalancing and passive ETF buying that lifts the entire market ecosystem.

### 5.2 Isolating True Sector Alpha

The true quantitative value of the World 3 Impulse Matrix is measuring the relative outperformance against the benchmark. By including the S&P 500 Index as a Reaction Asset, portfolio managers can isolate True Alpha above the market beta baseline. If a Nvidia shock historically drives the S&P 500 up 2.50% over a 5-day horizon, that 2.50% becomes the Market Beta baseline. If the same shock drives a peer technology stock up 4.00%, the engine has mathematically proven a 1.50% True Sector Alpha.

### 5.3 The Adobe-to-Nvidia Shockwave: Updated Case Study (2021–2026)

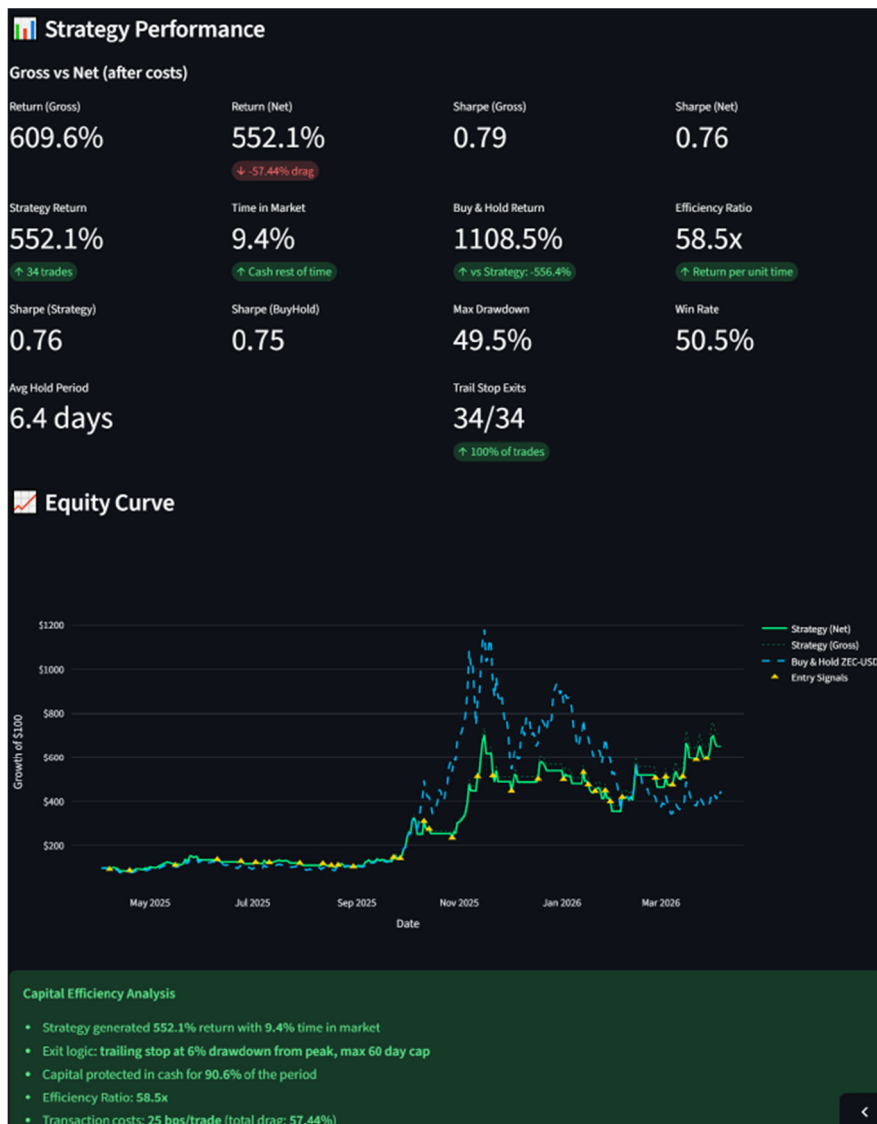
The flagship case study is the Adobe (ADBE) to Nvidia (NVDA) impulse. When Adobe experiences a Z-score shock, capital has historically rotated aggressively into Nvidia over the subsequent trading days. The strategy has been re-run against the full 2021–2026 dataset with transaction costs applied at the Low preset (10 bps round-trip).

| Parameter / Metric        | Value   |
|---------------------------|---|
| Strategy                  | ADBE Shock → NVDA Position  |
| Entry                     | T+1 after ADBE $ Z  \geq 1.5$ shock                               |
| Exit Logic                | Trailing Stop — 16% drawdown from peak                            |
| Max Hold Cap              | 20 trading days (case study setting; platform default is 45 days) |
| Z-Score Threshold         | $ Z  \geq 1.5$  |
| Backtest Period           | 2021 – 2026 (5 years)   |
| Transaction Costs Applied | 10 bps round-trip (Low preset)                                    |
| Total Trades              | 40  |
| Trail Stop Exits          | 4 / 40 (10% of trades)  |
| Strategy Return (Gross)   | +457.7%   |
| Strategy Return (Net)     | +405.2% (cost drag: -52.5%)                                       |
| Buy and Hold NVDA         | +3,251.5%   |
| Time in Market            | 37.0% (cash 63% of the time)                                      |

|                      |  |
|----------------------|--|
| Sharpe Ratio (Net)   | 0.78                                       |
| Sharpe Ratio (Gross) | 0.82                                       |
| Buy & Hold Sharpe    | 1.09                                       |
| Efficiency Ratio     | 11.0x (return per unit of market exposure) |
| Max Drawdown         | 51.1%                                      |
| Win Rate             | 53.3%                                      |
| Average Hold Period  | 19.1 days                                  |

The strategy maintains a net Sharpe of 0.78 after real-world transaction costs, confirming that the Adobe-to-Nvidia structural impulse relationship is robust to execution friction. The 37.0% time-in-market metric demonstrates the core capital efficiency thesis: the strategy captures meaningful returns while remaining in cash — or available for risk-free yield — for 63% of the backtest window.

[Figure 1: ADBE → NVDA Equity Curve (2021–2026) — Strategy Net vs Buy & Hold vs Entry Signals]



## 6. Exhibit D: Sector Rotation and Capital Friction (World 4)

World 4 is the newest addition to the Earth Scanner engine. It tracks the 11 GICS sector ETFs (XLK, XLF, XLE, XLV, XLI, XLP, XLY, XLB, XLU, XLRE) plus VTI as a total market benchmark, with full OHLCV history from 2000 to present (68,000+ daily bars, 25+ years).

While World 1 established that macro shocks create a rising tide across correlated assets, World 4 allows for the observation of Capital Friction — the measurable resistance and rotation that occurs when sector-specific shock energy dissipates into the broader economy.

### 6.1 Sector Rotation vs. Broad Contagion

Traditional finance theory suggests that during a technology-led volatility event, capital does not simply vanish — it rotates into Defensive or Value sectors like Financials or Consumer Staples. World 4 quantifies this rotation precisely. By running the Impulse Matrix with Technology (XLK) as the Driver Asset, the engine maps the Expected Alpha across all other sector ETFs, revealing which sectors act as capital recipients (positive alpha) and which experience contagion (negative alpha).

### 6.2 The Sector Impulse Hierarchy

The 25-year dataset allows the engine to test sector impulse relationships across multiple complete economic cycles including the dot-com bust (2000–2002), the financial crisis (2008–2009), the COVID crash (2020), and the post-pandemic normalization (2022–2024). Relationships that maintain high Confidence Scores across all of these regimes represent genuine structural market dynamics.

### 6.3 Tradable Takeaway: Sector Momentum Timing

By combining the Confidence Leaderboard with World 4 sector data, portfolio managers can identify the specific sectors that most reliably amplify or absorb shocks from major economic drivers. This transforms the Earth Scanner from a simple alert system into a sector rotation timing tool.

#### Key Advantage of World 4

Unlike individual stocks (World 3) where idiosyncratic earnings risk can overwhelm the cross-asset signal, sector ETFs smooth out individual company noise, making the impulse relationships more statistically reliable and persistent across time. The 25-year dataset spans multiple complete economic cycles.

## 7. The Confidence Leaderboard: Signal Quality Ranking

Not all cross-asset impulse relationships are created equal. Two relationships can have identical Expected Alpha but fundamentally different reliability profiles. The Confidence Leaderboard is the primary tool for distinguishing persistent structural anomalies from temporary statistical coincidences.

### 7.1 The Walk-Forward Stability Test

For a selected Driver Asset, the engine partitions the full historical shock database into chronological folds and evaluates two metrics:

- **Sign Rate:** What percentage of chronological folds produced Expected Alpha in the same direction? A 100% sign rate means directional consistency across every market regime tested.
- **Magnitude Retention:** How much of the training-period Expected Alpha survived into the test period? High retention means the edge is not decaying over time.

### 7.2 Nightly Pre-Computation

For computational efficiency at scale, the Confidence Leaderboard for default Driver Assets in each world is pre-computed nightly and stored in the conviction\_cache table in Supabase. This ensures instant load times for subscribers, while live walk-forward computation remains available for any custom Driver selection. Non-default driver selections are computed live and cached in the user session.

### 7.3 Interpreting the Leaderboard

| Score Range | Tier            | Interpretation   |
|-------------|-----------------|--|
| 65 - 100    | High Confidence | Persistent structural anomaly, tradeable across regimes      |
| 40 - 65     | Moderate        | Valid signal — verify with volume confirmation before sizing |
| 0 - 40      | Low / Unstable  | Likely regime-specific or statistical coincidence            |

## 8. The Nightly Automation Pipeline

Earth Scanner v3 operates on a fully automated nightly pipeline hosted on Railway, executing on a weekday schedule at market close (5:00 PM ET / 21:00 UTC). This pipeline runs a five-step sequence across all four Worlds, ensuring that every subscriber opens the platform each morning to a fully fresh, pre-computed dataset. A separate hourly cron job runs 24/7 for live intraday shock detection across crypto and open equity markets.

### Pipeline Sequence

- Step 1: Pre-flight freshness check — identify stale worlds before fetching
- Step 2: Price data fetch — yfinance (Worlds 1, 3, 4) + Binance (World 2)
- Step 3: Shock detection — adaptive Z-score windows per world
- Step 4: Conviction cache pre-computation — walk-forward leaderboards to Supabase
- Step 5: Distribution — X (Twitter) post + conditional Substack report

### 8.1 Multi-Source Price Fetching

World 1 (Macro) and Worlds 3 and 4 (Equities and Sectors) fetch daily OHLCV via yfinance. World 2 (Alt-Crypto) fetches via the Binance public API, which offers higher granularity and more complete altcoin coverage. All intraday data — 1-hour OHLCV bars for the most recent 30 days — is stored in a separate `prices_intraday` table, creating a hot/cold storage architecture that keeps read performance high while maintaining full historical depth.

### 8.2 Shock Detection at Scale

| World                | Symbols | Baseline | Z-Threshold    |
|----------------------|---------|----------|----------------|
| World 1: Macro       | 31      | 90 days  | $ Z  \geq 1.5$ |
| World 2: Alt-Crypto  | 154     | 15 days  | $ Z  \geq 2.0$ |
| World 3: S&P 500     | 503     | 90 days  | $ Z  \geq 1.5$ |
| World 4: Sector ETFs | 11      | 90 days  | $ Z  \geq 1.5$ |

Shocks are ranked by absolute Z-score across all worlds. If any asset breaches the extreme shock threshold ( $|Z| \geq 3.0$ ), an email alert is triggered independently of the scheduled distribution. The top three movers by absolute Z-score are selected for the daily X post.

### 8.3 Distribution: X and Substack

The pipeline composes a daily post for X using the top three shocks and the highest-confidence Expected Alpha signal from the Conviction Cache. Substack publication is conditional: a full shock report is posted on Fridays as a weekly recap, and additionally whenever an extreme shock ( $|Z| \geq 3.0$ ) is detected on any non-Friday trading day.

## 9. The Nightly Conviction Cache

Walk-forward stability analysis is computationally expensive. For a single Driver Asset in World 1 with 30 Reaction Assets and 4 chronological folds, the engine must execute over 240 individual event studies. At scale — 700+ assets across four worlds — running this analysis live on every user interaction would produce unacceptable load times.

The Conviction Cache solves this by pre-computing the full walk-forward leaderboard for each world's default Driver Assets nightly and storing the results as JSONB in Supabase. When a subscriber opens the Conviction tab and selects a default driver, the leaderboard renders instantly from cache.

### 9.1 Default Drivers Per World

| World                | Default Driver Assets (Pre-Computed Nightly)           |
|----------------------|--|
| World 1: Macro       | UUP (USD Index), GC=F (Gold), BTC-USD, ^GSPC (S&P 500) |
| World 2: Alt-Crypto  | BTC-USD, ETH-USD, SOL-USD                              |
| World 3: S&P 500     | AAPL, NVDA, MSFT, ^GSPC                                |
| World 4: Sector ETFs | XLK (Technology), XLF (Financials), XLE (Energy)       |

### 9.2 Cache Table Schema

| Field       | Description   |
|-------------|---|
| world       | World key (w_maj, w_alt, w_stocks, w_sectors)   |
| driver      | Driver asset ticker   |
| fwd_days    | Signal horizon (default: T+5)   |
| z_thresh    | Shock detection threshold (default: 1.5)  |
| computed_at | Timestamp of last nightly computation   |
| results     | JSONB array: ranked Reaction Assets with Expected Alpha, Confidence Score, Sign Rate, Mag Retention, N events |
| n_shocks    | Total historical shocks found for this driver at this threshold   |

## 10. The Forecast Tab: Confidence Bands and Outcome Distributions

The Forecast Tab in v3 surfaces the full statistical profile of historical post-shock outcomes, giving traders a complete picture of both the central tendency and the distribution of possible results — beyond the single Expected Alpha figure presented in earlier versions.

### 10.1 Pre-Shock Z-Score Warning

| Status       | Condition                            | Meaning   |
|--------------|--------------------------------------|---|
| SHOCK ACTIVE | $ Z  \geq \text{threshold}$          | Impulse event registered today — post-shock alpha analysis immediately relevant |
| APPROACHING  | $ Z  \geq 75\% \text{ of threshold}$ | Pressure building — monitor for threshold breach in next 1–2 sessions           |
| QUIET        | $ Z  < 75\% \text{ of threshold}$    | Normal market conditions — no actionable signal today                           |

### 10.2 Fan Charts with Confidence Bands

The fan chart shows the full distributional range of historical outcomes at each horizon from T+1 through T+7:

| Visual Element     | Definition   |
|--------------------|--|
| Cyan line          | Mean expected return (Expected Alpha from Section 2.2)           |
| Purple dotted line | Median return — less sensitive to fat-tail outliers              |
| Dark band          | 50% of outcomes — interquartile range (P25 to P75)               |
| Light band         | 80% of outcomes (P10 to P90) — captures most realistic scenarios |

Band width is itself an analytical signal. Narrow bands indicate a historically tight outcome distribution. Wide bands indicate high variance — the mean alpha may be attractive but the path is unpredictable. When the mean falls outside the 50% band, fat-tail outliers are skewing the average, and the median provides a more conservative sizing reference.

### 10.3 Outcome Distribution Histogram

For any selected forward horizon, the Forecast Tab renders a histogram of all historical post-shock returns, exposing characteristics that summary statistics cannot convey: return skew, bimodality, true win rate, and the full range of observed outcomes including worst and best historical results.

## 11. Capital Efficiency Backtester v2

The Portfolio Tab in v3 introduces three major upgrades: multi-driver entry logic, transaction cost modeling, and Monte Carlo statistical validation.

### 11.1 Multi-Driver Entry Logic

| Combination Mode | Entry Condition  |
|------------------|--|
| ANY (OR)         | Entry fires when any driver experiences a qualifying Z-score shock on the same day                     |
| ALL (AND)        | Entry fires only when every specified driver shocks on the same day — maximum conviction, fewer trades |
| MAJORITY         | Entry fires when more than 50% of the specified drivers shock on the same day                          |

### 11.2 Transaction Cost Modeling

| Cost Profile | Round-Trip Cost | Suitable For                                     |
|--------------|-----------------|--|
| Zero         | 0 bps           | Paper trading benchmarks only                    |
| Low          | 10 bps          | Large-cap equities, major crypto (BTC, ETH)      |
| Medium       | 25 bps          | Mid-cap stocks, liquid altcoins                  |
| High         | 55 bps          | Illiquid assets, exotic pairs                    |
| Custom       | User-defined    | Precise specification for institutional analysis |

### 11.3 Monte Carlo Statistical Validation

The Monte Carlo module runs 1,000 permutation simulations. In each simulation, actual shock dates are replaced with an equal number of randomly selected dates from the same historical window, producing a null distribution of returns from a random entry strategy with no structural edge.

The actual strategy return is compared against this null distribution. A return in the top 5% ( $p\text{-value} < 0.05$ ) provides strong statistical evidence that the edge is real rather than coincidental. This validation is particularly important for the ALL multi-driver mode, where a small sample size could otherwise produce misleadingly high returns from a handful of lucky trades.

## 12. Real-Time Hourly Alert Pipeline

While the nightly batch run provides the foundation of Earth Scanner's shock detection, markets don't wait for close-of-business. The Hourly Alert Pipeline extends the system into real time, scanning all four worlds every 60 minutes during live trading hours and updating both the website's Live Shock Scanner widget and the email alert system.

### 12.1 Architecture Overview

The hourly alert runner operates as a standalone Railway cron job, executing once per hour on a simple schedule. Unlike the nightly batch which processes the full analytical stack (forecasts, conviction grids, regime detection), the hourly runner is purpose-built for speed: it fetches the latest 1-hour OHLCV bar for each symbol, computes intraday Z-scores against the 90-day daily baseline, and writes the results to both the website database and the email pipeline.

| Component       | Details   |
|-----------------|---|
| Execution       | Railway cron, every 60 minutes                                |
| Data Source     | yfinance 1h OHLCV bars (2-day lookback for continuity)        |
| Baseline        | 90-day rolling daily returns (same as nightly)                |
| Storage         | Supabase PostgreSQL: prices_intraday_{world} + latest_shocks  |
| Website Update  | Top 5 movers by  Z  written to latest_shocks table every hour |
| Alert Threshold | Z  ≥ 3.0 (extreme shocks only) using time-scaled Z-scores     |
| Deduplication   | 24-hour cooldown per symbol/world/alert type                  |
| Cap             | Maximum 5 alerts per email dispatch                           |

### 12.2 Intraday Z-Score Computation

The core innovation of the hourly pipeline is the intraday Z-score, which compares the current price movement against the 90-day daily return distribution. This uses the same statistical framework as the nightly run but applies it to partial-day data.

#### The Time-Scaling Problem

A naive Z-score computed at 10:00 AM would compare a 30-minute return against a full-day return distribution, systematically underestimating true shock magnitude. Earth Scanner solves this with a square-root-of-time scaling factor:

#### Intraday Z-Score Formula

$z_{\text{raw}} = (\text{intraday\_return} - \mu) / \sigma$  where  $\mu$  and  $\sigma$  come from 90-day daily returns

$z_{\text{scaled}} = z_{\text{raw}} \times \sqrt{(24 / \text{hours\_passed})}$  for alert threshold decisions

The raw Z is displayed on the website and in emails. The scaled Z is used only for the  $|Z| \geq 3.0$  alert threshold, ensuring that an early-morning move of the same magnitude as an end-of-day move receives equal treatment.

## Two Z-Scores, Two Purposes

The system maintains two Z-score variants for each symbol. The raw Z-score ( $z_{\text{raw}}$ ) measures the actual intraday return against the daily baseline without adjustment and is used for all display purposes: the website's Live Shock Scanner widget, email alert text, and the app's shock panels. The time-scaled Z-score ( $z_{\text{scaled}}$ ) applies the square-root-of-time adjustment and is used exclusively for alert threshold decisions. This separation ensures that users always see the true magnitude of the move, while the alerting system compensates for partial-day observation windows.

## 12.3 Market Hours Awareness

The pipeline respects the fundamental difference between 24/7 crypto markets and traditional equity markets. Crypto worlds (W2: Altcoin Pairs) fetch live intraday data at every hourly tick regardless of time. Equity worlds (W3: S&P 500 Stocks, W4: Sector ETFs) only fetch live bars during US market hours (9:30 AM – 4:00 PM Eastern, Monday through Friday). Outside market hours, equity worlds use the last known daily close for their Z-score baseline. The macro world (W1: Major Pairs) includes both crypto and traditional assets and runs on the crypto schedule since at least some of its components trade continuously.

## 12.4 Live Website Widget

Every hourly run writes the top 5 movers (by absolute Z-score, filtered to  $|Z| \geq 1.5$ ) to the latest\_shocks table in Supabase. The website's Live Shock Scanner widget fetches this table via a lightweight JavaScript call and renders the current market movers with their Z-scores, percentage changes, direction indicators, severity labels, and the current regime context.

| Severity | $ Z $ Range | Interpretation                     |
|----------|-------------|------------------------------------|
| Moderate | 1.5 – 1.99  | Notable deviation from baseline    |
| Severe   | 2.0 – 2.99  | Significant market stress          |
| Extreme  | $\geq 3.0$  | Rare event, potential regime shift |

## 12.5 Conviction-Enriched Alerts

When an extreme shock triggers an email alert, the system enriches it with conviction data from the pre-computed conviction cache. For each shocked driver symbol, it queries the most recent 5-day forward conviction results and attaches the best signal that meets minimum quality thresholds: confidence score  $\geq 40$  and expected alpha  $\geq 1.5\%$ . This transforms a simple shock notification into an actionable trading signal, telling the subscriber not just that a shock occurred, but which reaction asset has historically benefited and with what expected magnitude.

## 12.6 Suspect Tick Filtering

To guard against data quality issues (stale prices, exchange glitches, stock splits), the alert system applies per-world return caps before promoting any shock to alert status. Crypto (W2) caps at 50%, equities (W3) at 20%, and macro/sectors (W1/W4) at 15%. Any tick exceeding these thresholds is logged as suspect and excluded from email alerts, though it may still appear on the website widget for informational purposes.

## 13. Z-Decay Exit Strategy

Traditional backtesting exits rely on fixed holding periods or trailing stops. Earth Scanner's Capital Efficiency Backtester introduces a third, signal-aware exit: Z-Decay. This mechanism monitors the driver's Z-score after entry and exits when the statistical shock that triggered the trade has dissipated, regardless of whether the price target or trailing stop has been reached.

### 13.1 The Problem with Fixed Exits

A shock-driven entry thesis has a natural expiration. If the driver's Z-score was  $-3.2$  at entry (a severe bearish shock), and three days later it has normalized to  $-0.4$ , the original thesis is spent. Holding the position through a 45-day max hold period wastes capital on a dead signal. Conversely, a trailing stop exits on price action alone and may trigger during normal volatility even while the shock thesis remains intact.

### 13.2 How Z-Decay Works

The Z-Decay exit adds a third exit condition alongside the trailing stop and max hold period. On each day of an open position, the system checks the driver's current absolute Z-score. If it falls below the configurable Z-decay threshold, the position is closed with the exit reason logged as the specific threshold value.

#### Z-Decay Exit Rule

IF  $|current\_z\_score| < z\_decay\_threshold$  THEN exit position

The Z-decay threshold is user-configurable in the Portfolio tab. Typical values range from 0.5 to 1.0. A threshold of 0.5 means: exit when the shock has faded to within half a standard deviation of normal.

### 13.3 Three-Layer Exit Hierarchy

Each open position is evaluated against three exit conditions on every trading day, checked in the following order:

| Priority | Exit Type     | Logic   |
|----------|---------------|---|
| 1        | Trailing Stop | Peak-to-current drawdown exceeds <code>trail_pct</code> (default 10–20% by world) |
| 2        | Z-Decay       | Driver's $ Z $ drops below <code>z_decay_threshold</code> (signal exhaustion)     |
| 3        | Max Hold      | Position has been open for <code>max_hold</code> days (default 45 days)           |

The trailing stop provides downside protection against adverse price moves. Z-Decay provides thesis protection by recognizing when the entry signal has expired. Max hold provides time protection as a hard backstop. The trade log records which exit condition triggered for each position, enabling users to analyze exit effectiveness across their backtest.

## 13.4 Default Parameters by World

| World              | Trail % | Max Hold | Z Thresh | Cost Tier |
|--------------------|---------|----------|----------|-----------|
| W1: Major Pairs    | 10%     | 45 days  | 1.2      | Low       |
| W2: Altcoin Pairs  | 20%     | 45 days  | 1.5      | Low       |
| W3: S&P 500 Stocks | 10%     | 45 days  | 1.2      | Low       |
| W4: Sector ETFs    | 10%     | 45 days  | 1.2      | Low       |

The wider 20% trailing stop for crypto reflects the higher intraday volatility of altcoins, which can swing 15–20% within a single rally leg. The higher Z threshold (1.5 vs 1.2) filters out the noise inherent in crypto markets, requiring a stronger initial signal before entering.

## 14. Shock Clustering and Named Macro Events

Individual asset shocks rarely occur in isolation. When gold rallies, treasuries move, and equities sell off simultaneously, these aren't three independent events — they're one macro narrative: a flight to quality. Earth Scanner's shock clustering engine automatically detects these correlated shock groups and classifies them into named macro events, transforming a list of ticker-level signals into a coherent market story.

### 14.1 Correlation-Based Clustering

The clustering algorithm operates in three steps. First, it computes the rolling pairwise correlation matrix among all currently-shocked assets using 60 days of return history. Second, it builds a correlation graph where each shocked asset is a node and an edge connects any two assets with absolute correlation  $\geq 0.35$ . Third, it extracts connected components using breadth-first search, producing clusters of correlated shocks and leaving uncorrelated shocks as singletons.

#### Clustering Parameters

Correlation Lookback: 60 trading days

Correlation Threshold:  $|r| \geq 0.35$

Method: BFS connected components on the correlation adjacency graph

These defaults balance sensitivity (catching genuine macro linkages) with specificity (avoiding spurious connections from transient noise).

### 14.2 Asset-Class Taxonomy

Every symbol in the Earth Scanner universe is mapped to an asset class for event classification. The taxonomy covers 14 classes across all four worlds: dollar, rates, equity\_index, volatility, energy, precious\_metals, industrial\_metals, agriculture, crypto\_major, crypto\_alt, fx, fx\_em, equity (for S&P 500 stocks), and equity sector-specific classes derived from metadata. This classification enables the pattern matching that transforms raw clusters into named events.

### 14.3 Named Macro Event Templates

Each cluster is matched against a library of macro event templates based on which asset classes are present and the dominant direction (bullish, bearish, or mixed). The system checks templates in priority order and assigns the first match.

| Event Name              | Required Classes               | Interpretation                             |
|-------------------------|--------------------------------|--|
| Flight to Quality       | precious_metals + equity_index | Safe havens rallying, risk assets falling  |
| Risk-On Surge           | equity_index + crypto (up)     | Broad risk appetite across asset classes   |
| USD Liquidity Event     | dollar + rates + equity_index  | Macro liquidity repricing across markets   |
| Dollar Repricing        | dollar + fx                    | USD move rippling through currency pairs   |
| Energy Shock            | energy                         | Supply/demand or geopolitical energy event |
| Volatility Regime Shift | volatility + equity_index      | VIX signaling a new volatility regime      |
| Crypto Cascade          | crypto_major + crypto_alt      | Bitcoin shock propagating through altcoins |
| Rates Shock             | rates                          | Treasury market stress or policy repricing |
| Commodity Repricing     | energy + precious_metals       | Broad commodity inflation/supply signal    |
| EM Currency Stress      | fx_em                          | Emerging market currencies under pressure  |
| Sector Rotation         | equity (sector-specific)       | Movement concentrated in specific sectors  |

### 14.4 Cluster Scoring

Each cluster is scored on two dimensions that together determine its prominence in the dashboard. Breadth measures how many assets are participating (wider = more significant macro event). Intensity measures the mean absolute Z-score across cluster members (stronger = more extreme statistical deviation). The composite score uses total Z (sum of absolute Z-scores across all members), which naturally weights both breadth and intensity. Clusters are ranked by total Z descending, so a 4-asset cluster averaging  $|Z| = 2.5$  (total Z = 10.0) outranks a 2-asset cluster averaging  $|Z| = 4.0$  (total Z = 8.0) — reflecting that the broader event carries more systemic significance.

### 14.5 Within-Cluster Correlation

Each cluster reports its average within-cluster correlation as a quality metric. A “Flight to Quality” cluster with  $avg\_correlation = 0.72$  means the gold, treasury, and equity moves are tightly synchronized, reinforcing the narrative. A cluster with  $avg\_correlation = 0.38$  (barely above the 0.35 threshold) suggests a looser connection that may warrant skepticism. This metric helps users distinguish between strong macro narratives and coincidental co-movement.

## 14.6 Integration with the Dashboard

Shock clusters surface throughout the Earth Scanner interface. In the Forecast tab, clusters provide context for individual asset shocks by showing which broader narrative they belong to. In the Macro tab, the regime detection system uses cluster breadth as an input signal: a 6-asset “Flight to Quality” cluster is strong evidence of a risk-off regime. In the hourly alert emails, cluster context helps subscribers understand not just that gold spiked, but that it spiked as part of a coordinated safe-haven rotation. The goal is always the same: transforming statistical signals into actionable market narratives.

### Fallback Classification

When no template matches a cluster’s asset-class composition, the system applies size-based fallbacks: clusters with 5+ assets become “Broad Market Event,” 2–4 assets become “Correlated Shock Pair,” and singletons become “Isolated Shock.” These ensure every detected shock receives a classification, even for novel combinations.

## 15. Platform Architecture

Earth Scanner v3 is deployed as two distinct Streamlit applications sharing a common engine, data pipeline, and Supabase PostgreSQL backend.

### 15.1 Two-Tier Architecture

| Component        | Free (World 1)                      | Pro (All 4 Worlds)                      |
|------------------|-------------------------------------|---|
| Authentication   | None                                | Google OAuth via Streamlit              |
| Access Gate      | Terms of Service acceptance         | Active Stripe subscription + TOS        |
| Worlds Available | World 1 only                        | All 4 worlds (700+ assets)              |
| Tab Layout       | 4-tab simplified or 6-tab advanced  | 6-tab advanced + admin Data/Diagnostics |
| Beginner Mode    | Available — plain-English summaries | Not shown                               |

### 15.2 Terms of Service Gate

All access to Earth Scanner — both free and Pro — requires explicit Terms of Service acceptance before any data is displayed. The TOS checkbox is a mandatory gate on the login screen for the Pro application and on first visit for the free application. No scan, no data load, and no Stripe redirect is accessible until the user has affirmatively accepted the terms.

### 15.3 Subscription and Access Verification

Pro subscriptions are processed through Stripe at \$99/month (founding members locked at \$79/month). A Vercel serverless webhook listens for Stripe subscription events and writes confirmed subscriber email addresses to the subscribers table in Supabase. On every Pro app load, the access layer queries this table in real-time — subscription cancellations take effect on the next platform load with no manual intervention required.

## 15.4 Database Architecture

| Table                    | Description  |
|--------------------------|--|
| prices_daily__{world}    | Full OHLCV history at daily granularity — primary read source for all analysis |
| prices_intraday__{world} | 1-hour OHLCV bars, rolling 90-day retention — used for volume confirmation     |

*End of Document*

© 2026 7th Dimension Markets. All rights reserved.