CURIOSITYBITS LIVING LAB BOSTON · STUDY #03 PRE-REGISTRATION & STUDY DESIGN CARD (V2.1)

Pre-Registration / Study Design Card — Study #03: Tech Community Skepticism vs. Wall Street Hype (`SKILL-012`)

1. Research Question & Theoretical Anchor

- **Plain-Language Question:** When publicly traded tech companies announce major product launches, AI models, cloud platforms, or strategic pivots (2023–2026), Wall Street often reacts immediately to the headline press release on Day 1. Meanwhile, software engineers and practitioners on **Hacker News (`news.ycombinator.com`)** debate within hours whether the technology actually works, is overpriced, or is vaporware. Does the gap between Wall Street's Day-1 price reaction and grassroots engineer skepticism on Hacker News predict 30-day post-event abnormal stock returns (CAR[+1, +30])?
- **Theoretical Anchors (`Field Guide`):** 1. **Wisdom of Crowds vs. Information Cascades (Surowiecki, 2004; Bikhchandani, Hirshleifer, & Welch, 1992):** Specialized practitioner communities possess distributed technical knowledge that may take weeks to diffuse into institutional equity research. 2. **Limited Investor Attention & Complexity (Hirshleifer, Lim, & Teoh, 2009):** Equity investors react quickly to headline announcements on Day 1 (`Day1_Return`) but underreact to granular technical defects or developer switching friction surfaced in practitioner forums. 3. **Narrative Economics & Cheap Talk (Shiller, 2017; Crawford & Sobel, 1982):** Press-release claims are low-cost signaling ("cheap talk"), whereas hands-on practitioner benchmarks and failure reports act as costly verification.

2. Q1 — Exact Operational Measurement

- **Open-Access Data Sources (`knowledge/open_access_data_connector.py`):** 1. **Hacker News Algolia REST API (`hn.algolia.com/api/v1/search` & `/items/{id}`):** Live retrieval of top Hacker News stories and practitioner comment threads discussing major public-company tech/AI product launches, earnings/pricing shifts, and platform announcements across 2023–2026. 2. **Wikimedia Daily Pageviews API (`wikimedia.org/api/rest_v1/metrics/pageviews/`):** Public attention spike ratio around the announcement date. 3. **Yahoo Finance (`yfinance`):** Daily OHLCV stock prices for each firm and `SPY` benchmark.
- **Unit of Analysis & Sample Attrition Rules (`SKILL-006` & `SKILL-011`):**
- Target initial candidate pool: $N_{in} \geq 45$ major public-company product/platform/earnings events discussed on Hacker News (2023–2026).

- Mandatory non-overlapping filter (`SKILL-006`): Enforce a strict ≥ 35 calendar-day cooldown between events for the same ticker so 30-day post-event return windows never overlap (N_{out} ≥ 35 strictly non-overlapping events across ≥ 20 distinct public tech firms).
- **Primary Independent Variables (IVs):**
- `HN_Technical_Skepticism` (S_{HN} \in [0, 1]): Multi-model consensus score measuring technical doubt, vaporware/moat skepticism, pricing pushback, and reported engineering flaws in Hacker News titles and top practitioner comments.
- `HN_Practitioner_Verification` (V_{HN} \in [0, 1]): Multi-model consensus score measuring hands-on developer praise, benchmark verification, and real production adoption in Hacker News discourse.
- `Hype_Skepticism_Divergence` (D_{HS}): Standardized divergence between Wall Street's Day-1 price reaction (`Day1_Return`) and Hacker News technical reception (S_{HN} - V_{HN}), capturing events where Wall Street bids up a stock on Day 1 while engineers pan the product (or vice versa).
- **Primary Dependent Variable (DV):**
- 30-day market-adjusted Cumulative Abnormal Return:

> *CAR_{i,[+1,+30]} = R_{i,[+1,+30]} - R_{SPY,[+1,+30]}*

3. Q2 — Text Annotation & Open-Weight / JEV Validation Plan

- **Primary Lead Model:** `Gemini 4 Argon` (`temp = 0.0`)
- **Open-Weight & JEV Cross-Annotators (`SKILL-012`):** 1. `deepseek/deepseek-v4.1-flash` (Open-Weight) 2. `qwen/qwen3.8-flash` (Open-Weight) 3. `typesafe/jev-router` (JEV Joint Evaluation Decision Model) 4. `local/hn-structural-lexical-v1` (Deterministic technical-lexicon baseline)
- **Target Per-Construct Reliability (`SKILL-011` / `SKILL-012`):** Intraclass Correlation ICC(2,k) ≥ 0.75 and quadratic-weighted Cohen's κ ≥ 0.70 reported separately for `HN_Technical_Skepticism` (S_{HN}), `HN_Practitioner_Verification` (V_{HN}), and `Hype_Skepticism_Divergence` (D_{HS}) in `projects/hn-skepticism/data/annotation_validation.json`.

4. Q3 — Mandatory Confounders & Falsification Criteria

- **Baseline Financial & Attention Confounders Controlled (`SKILL-005`):** 1. `Day1_Return`: Immediate Day-1 stock return ([0, +1]) to absorb the headline announcement shock. 2. `Pre20D_Mom`: Pre-event 20-day cumulative stock return ([-20, -1]). 3. `Pre20D_Vol`: Pre-event 20-day annualized realized volatility. 4.

`Log_HN_Engagement`: Log of Hacker News story points + comment count (controlling for raw discussion volume). 5. `Sector_FE`: Sub-sector fixed effects (Semiconductors, Cloud/SaaS, Consumer Platforms, Enterprise AI) + Firm-Clustered Standard Errors (`SKILL-006`).

- **Pre-Stated Falsification Criterion (`SKILL-012`):**
- If neither `HN_Technical_Skepticism` (S_{HN} - V_{HN}) nor `Hype_Skepticism_Divergence` (D_{HS}) achieves $|t| \geq 1.96$ ($p < 0.05$) in the multivariate OLS regression after controlling for `Day1_Return`, `Pre20D_Mom`, `Pre20D_Vol`, `Log_HN_Engagement`, and `Sector_FE` on the non-overlapping panel, we commit to reporting a **null or partial finding** transparently on the lab website without p-hacking or changing the specification ex-post.

---

5. v2.0 Sample Expansion & Switching-Cost Boundary Addendum (`SKILL-008`, `SKILL-011`)

- **Why `v1.0` (N=43, M=20 firms) Was Archived (`PAPER_MANUSCRIPT_V1_ARCHIVED.md`):**
- In `v1.0`, the initial 50-query run yielded N=43 non-overlapping events ($34 front-page threads ≥ 25 points across 20 firms), including 9 low-visibility posts (<25 points) and 2 off-topic Algolia matches (`IBM #36446180` and `PLTR #36894613`) flagged by our independent `critical-peer-reviewer` subagent (`INDEPENDENT<sub>PEER</sub><sub>REVIEW</sub><sub>REPORT</sub>.md`). With only 20 clusters, firm-clustered standard errors were ± 7.04% (β̂ = -3.83%, p = 0.586$).
- **`v2.0` Large-Sample Expansion Protocol (N ≥ 150+ Non-Overlapping Front-Page Threads Across $50+$ Firms, 2022–2026):** 1. **Strict Front-Page & Title-Verified Inclusion Gate:** Every thread must have ≥ 40 Hacker News points, ≥ 20 comments, and an explicit company/product keyword match in the story title (`hn.algolia.com`), eliminating low-visibility noise and off-topic matches. 2. **Expanded Firm Universe (M ≥ 50 Public Tech Firms):** Covers Semiconductors, Enterprise Software & Cloud, Developer Tools/Infrastructure, and Consumer Platforms across 2022–2026 with a strict ≥ 35-calendar-day same-ticker cooldown (`SKILL-006`). 3. **Ex-Ante Switching-Cost / Lock-In Indicator (`Low_LockIn_Regime`):** Tests whether developer skepticism interacts with customer switching costs (`Low_LockIn_Regime = 1` for developer tools, consumer hardware/apps, open-source/cloud APIs, and chips vs. `0` for high-lock-in legacy enterprise ERP/virtualization/suite monopolies `ORCL`, `IBM`, `SAP`, `AVGO`, `CSCO`, `WDAY`, `NOW`, `INTU`, `ADBE`, `CRM`, `MSFT`).

---

6. Edition `v2.1` Sister-Lab Audit (`EXT-AUDIT-034`) & Pre-Registration Commit Addendum (`SKILL-019`, `SKILL-023`, `SKILL-024`, `SKILL-025`, `SKILL-027`)

- **Pre-Registration Commit Provenance Audit (`SKILL-024`):**

- Historical git log inspection (`experiments/GIT_COMMIT_HISTORY_AUDIT.txt`) confirms that in Study #03's initial commit (`dfb9353fdde0bd80143abc50e95c62a4b22b0e8a`, `2026-10-29 19:51:46 +0000`), `STUDY_DESIGN_CARD.md` (archived in `STUDY_DESIGN_CARD_V1_PRE_DATA_LOCKED.md`, SHA-256 `b002d2eefaf6b0d5292c43fbde0defadc746cb68a5c7ab6498c711bd3b19dffe`) was committed in the same git commit as the `v1.0` dataset (`N = 43`), and in commit `7f25c87395263f09524d06fe058b5d2b5180e9d0` (`2026-10-29 21:43:38 +0000`), Section 5 (`v2.0` addendum, archived in `STUDY_DESIGN_CARD_V2_0_ARCHIVED.md`, SHA-256 `c3c3c47dc80e7b049d2bfe9d16d1bfe70cb9203741387d70b527ca3c40ab5821`) was committed alongside the `v2.0` dataset (`N = 202`). Unlike Study #05 (where commit `2d8e8d1` locked the design card prior to data collection), Study #03's card and data were committed together; we therefore disclose this explicitly and treat `Model 3` (`Low_LockIn_Regime`) and the Section 4.3 event-archetype mechanism regressions as **exploratory / post-hoc boundary analyses**.

- **Explicit Testing of Pre-Registered `Hype_Skepticism_Divergence` (D_{HS} / `Hype_Skepticism_Trap`):**

- In `Edition v2.1`, we formally estimate and report **Model 4** testing `Hype_Skepticism_Trap` (D_{HS} = $Net_{Skepticism}$ + 0.04 × $Day1_{Return}$ × $\max(0, Net_{Skepticism})$) in both bivariate OLS ($\hat{\beta}$ = +1.1413%, SE = 3.5260%, t = +0.324, p = 0.74617, R^2 = 0.0004) and full Model 2 confounder-controlled OLS ($\hat{\beta}$ = -1.5087%, SE = 3.8843%, t = -0.388, p = 0.69771 ≈ 0.6977, R^2 = 0.0492), confirming that both pre-registered primary IVs (`Net_HN_Skepticism` and `Hype_Skepticism_Divergence`) satisfy the pre-stated falsification boundary ($|t| < 1.96$).

- **Model 5 Non-Overlapping Day-1 Drift Diagnostic & Exploratory Archetype Multiple-Testing Correction:**

- In **Model 5** over strictly non-overlapping (t+1, t+30] post-event returns (excluding shared trading session t+1), the apparent positive slope on `Day1_Return` (+0.4531, p = 0.00964 in Model 2 over (t_0, t+30]) drops to **-0.0053 (p = 0.97425; bivariate +0.0027, p = 0.98645)**, demonstrating that the Day-1 drift is a mechanical one-day window-overlap artifact.

- Each of the N = 202 threads is deterministically coded from its HN story title into four mutually exclusive archetypes (`Layoff_RTO` n = 12, `Product_Security_Incident` n = 15, `Antitrust_MA_Legal` n = 16, and `Other` n = 159, alongside a broader keyword sensitivity coding of n = 18 layoff/RTO and n = 23 product/security/pricing events). This archetype analysis is **exploratory**: the pooled layoff interaction (+55.0046%, unadjusted 999-shuffle label-permutation p = 0.0290) is **not statistically significant after a four-archetype Bonferroni correction (p = 0.1160 ≈ 0.116)** and attenuates to +11.6140% (p = 0.3910) across all $18 broad layoff keyword matches (`Layoff<sub>RTO</sub>` n = 12 subsample slope +34.7670%, t(7)\ p = 0.0162, HC3\ p = 0.0882$).

- **Post-Study Out-of-Sample Evaluation (`2026-08-08` to `2026-09-25`, N = 28 Events Across M = 28 Firms):**

- Evaluates $28 fresh front-page Hacker News threads (`points >= 40, num<sub>comments</sub> >= 20`) occurring after the `2026-08-08` end of the `v2.0` sample with complete 30-trading-day `yfinance` return windows (`data/hn<sub>skepticism</sub><sub>oos</sub><sub>2026</sub><sub>08</sub><sub>09</sub>.json`), reporting both SG `EXT-AUDIT-034` 's lexical rubric (r = -0.062, p = 0.755, 95% CI [-0.43, +0.32]) and Boston's `score<sub>lexical</sub><sub>hn</sub>` rubric (r = -0.2403, p = 0.2181, 95% CI [-0.562, +0.146]) alongside explicit minimum detectable effect size bounds (MDES r = 0.508 at 80% power, α = 0.05$).