

A Multi-Agent Deep Framework for Indonesian Single-Stock Trading with an ESG Factor

Evidence from **ANTM.JK** during the 2025 Gag Nikel
mining controversy.

Dingyue Kan · Yichong Bai — The Hong Kong Polytechnic University

Jerry Cao — The Hang Seng University of Hong Kong, Research Centre for ESG

ANTM.JK · LQ45 · IDX
CTDE · LLM · ESG GATING

Disclaimer

Information presented in the material is not an offer to sell or solicitation for an offer to buy any financial products and services and it should not be considered as investment advice. The parties involved in preparing the content of this presentation material are not liable financially in any way as a result of using information documented by any individuals or corporations. The Research Centre for ESG and the research team involved accept no responsibility or liability for the accuracy and completeness of such information presented in the event.

Abstract

This work proposes a new paradigm of algorithm trading system with AI agent utilizing ESG news and showcases how the system works effectively in the emerging market. ESG news is effectively traded with AI agent to generate superior performance. The research highlights the importance of ESG news. .

Three contributions, one case study

01

First ESG factor on the IDX

LLM-generated controversy signals drive single-stock trading inside a CTDE multi-agent architecture.

02

Modular & auditable

The ESG Agent never touches price data; the Trading Agent never sees raw news — separation built for financial-AI governance.

03

A rigorous case study

A 3,825-parameter grid and seven ablations on ANTM.JK across the 2025 Gag Nickel controversy.

SECTION 01

Motivation & Setting

Why ESG in emerging markets is a daily-frequency problem — and what the IDX does to a trader.

ESG risk arrives as news, not as filings

Annual sustainability reports and static ratings are **structurally too slow** for a resource firm. ANTM's controversies — permitting disputes, ecological allegations, regulatory suspensions — **erupt within days**.

They diffuse through high-context, unstructured **news text**, not structured filings. The market holds the raw information but lacks the semantic capacity to quantify it in real time.

THE INFORMATION ASYMMETRY

Raw public information is everywhere.
Real-time semantic interpretation is not.

Four constraints shape every feasible strategy

01 No naked short selling

Indonesia's intraday short programme was non-operational through 2025 — a long-only baseline is mandated.

03 Retail "BPJS trading"

A strong day-trading tradition encourages flattening positions by the close.

02 15% auto-rejection (ARB)

Since April 2025, sell orders are rejected at the circuit-breaker limit, capping exits during extreme down moves.

04 ~60 bps round-trip cost

Roughly six times US institutional rates, in a centralized book with no wholesalers absorbing flow.

Is the IDX efficient with respect to ESG news?

Under Fama's (1970) **semi-strong efficient market hypothesis**, public information is impounded instantly. If that held here, ESG-conditioned trading could not beat the 60 bps cost.

We treat this as a hypothesis and build a **temporally shuffled control** — the same signals, randomly reordered, stripped of their informational content.

THE PROBE

- Excess returns survive **only under correct order** → predictability is genuine; semi-strong form **violated**.
- Returns **vanish under shuffle** → the market efficiently digested the public news.

— VALIDATION TARGET

ANTM.JK — a blue-chip nickel miner on the LQ45

LQ45

BENCHMARK INDEX

The 45 most liquid, highly capitalized stocks on the IDX.

~60_{bps}

ROUND-TRIP COST ABSORBED

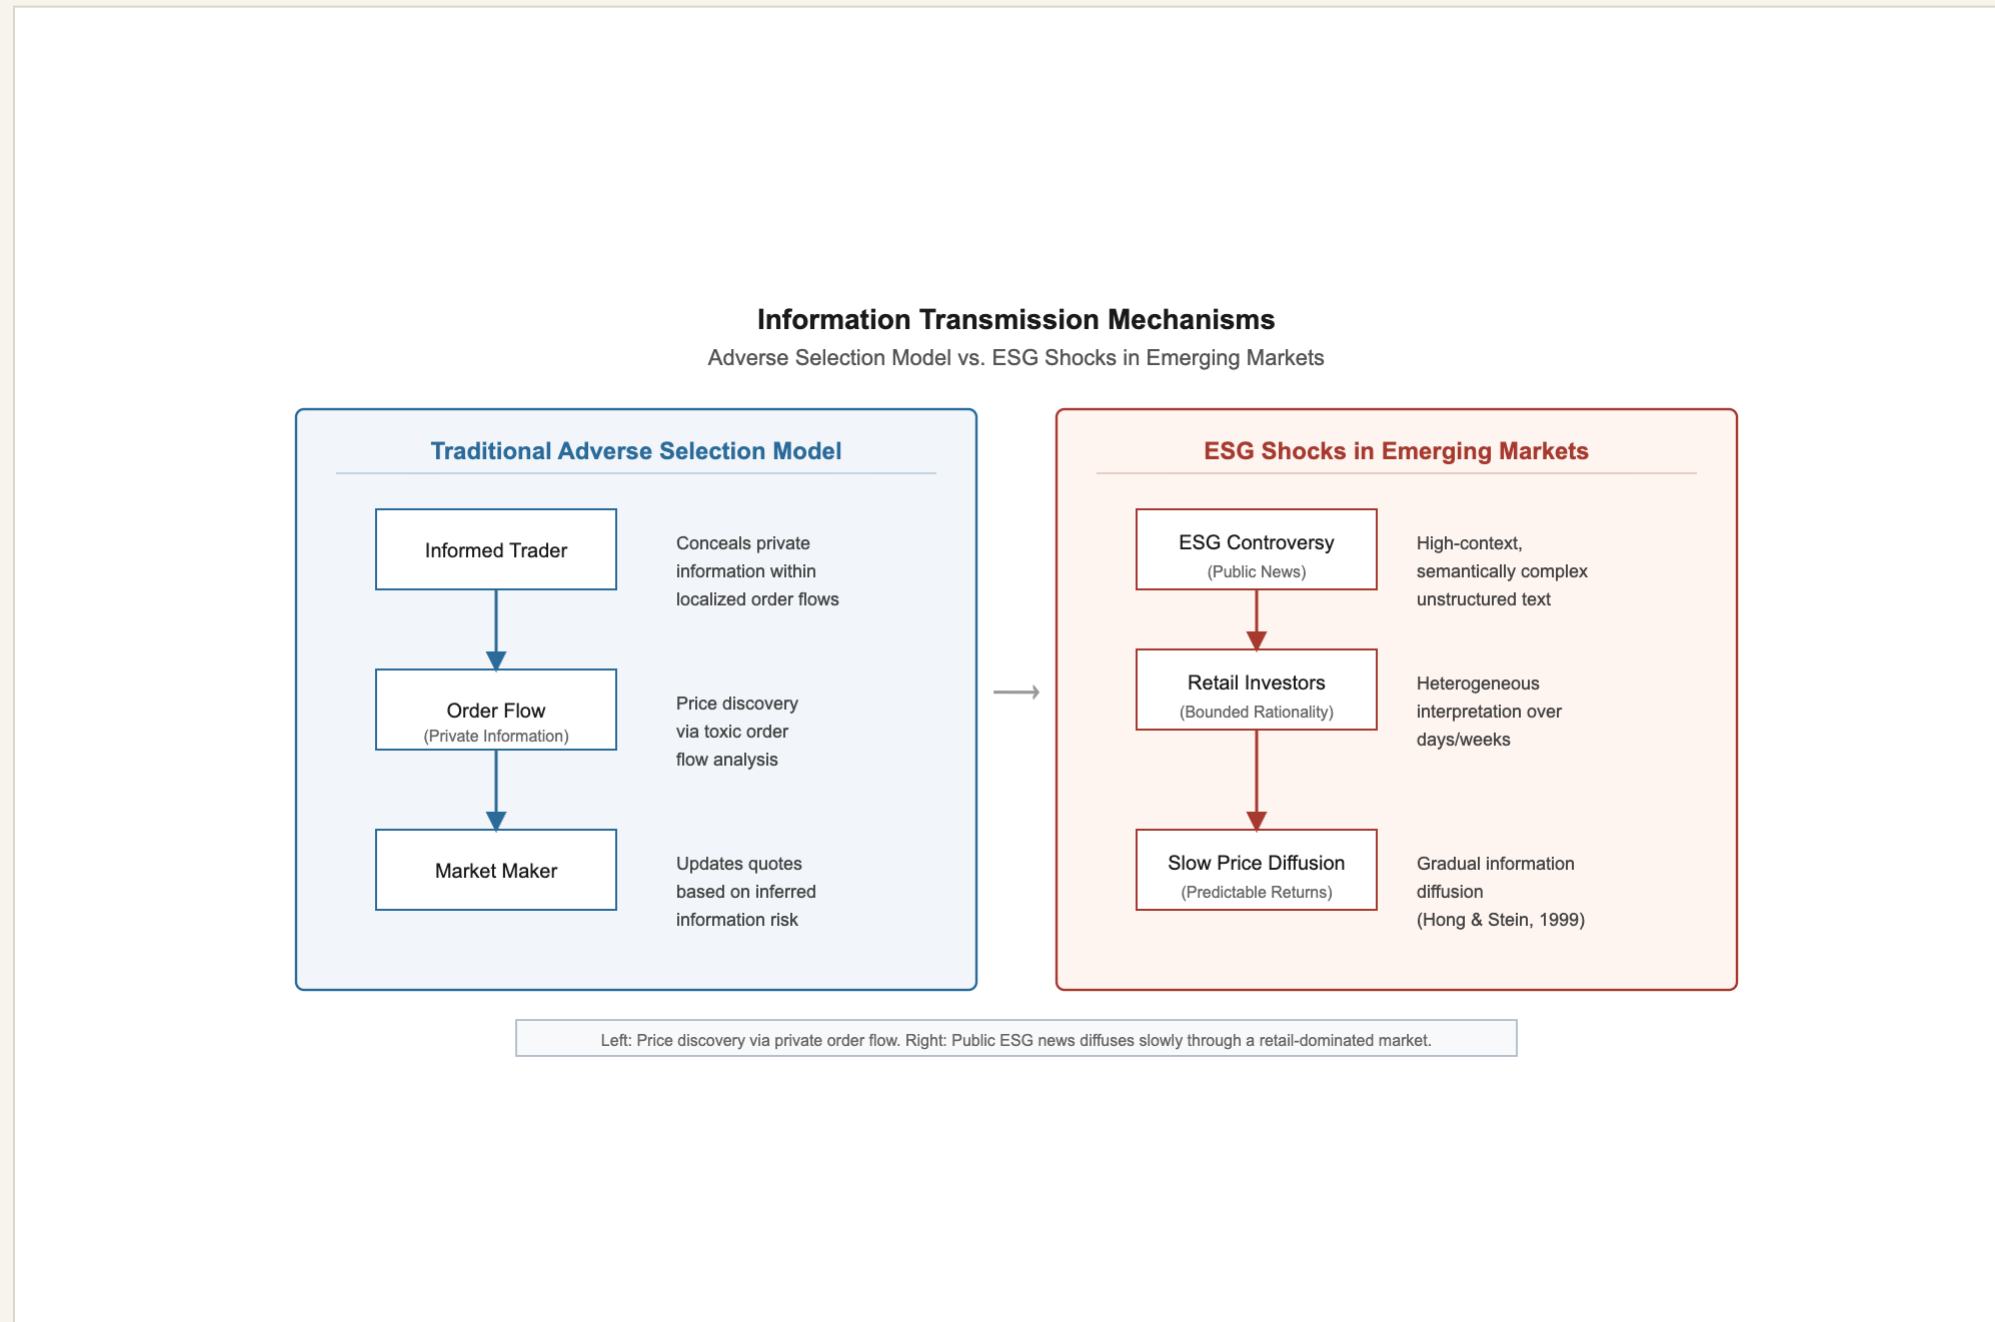
Liquidity provides depth to mitigate execution friction.

SOE

STATE-OWNED MINER

Directly exposed to environmental & governance controversy.

Microstructure models miss ESG shocks



Adverse-selection models assume price discovery through **private information hidden in order flow** . ESG controversy is the opposite — **public, high-context, semantically complex** .

Its impact diffuses through the **slow, heterogeneous interpretation of text** by a retail base — the gradual diffusion of Hong & Stein (1999).

A 30-day half-life signal is built for daily diffusion , not for microseconds.

SECTION 02

A Multi-Agent CTDE Framework

Three agents under a centralized-training,
decentralized-execution blueprint — instantiating the
Agentic AI paradigm.

Three agents, cleanly separated

● ESG AGENT

Perception

Crawls daily news, feeds source-backed articles to an LLM, and outputs a controversy intensity ξ_t . **Never sees price.**

● TRADING AGENT

Execution

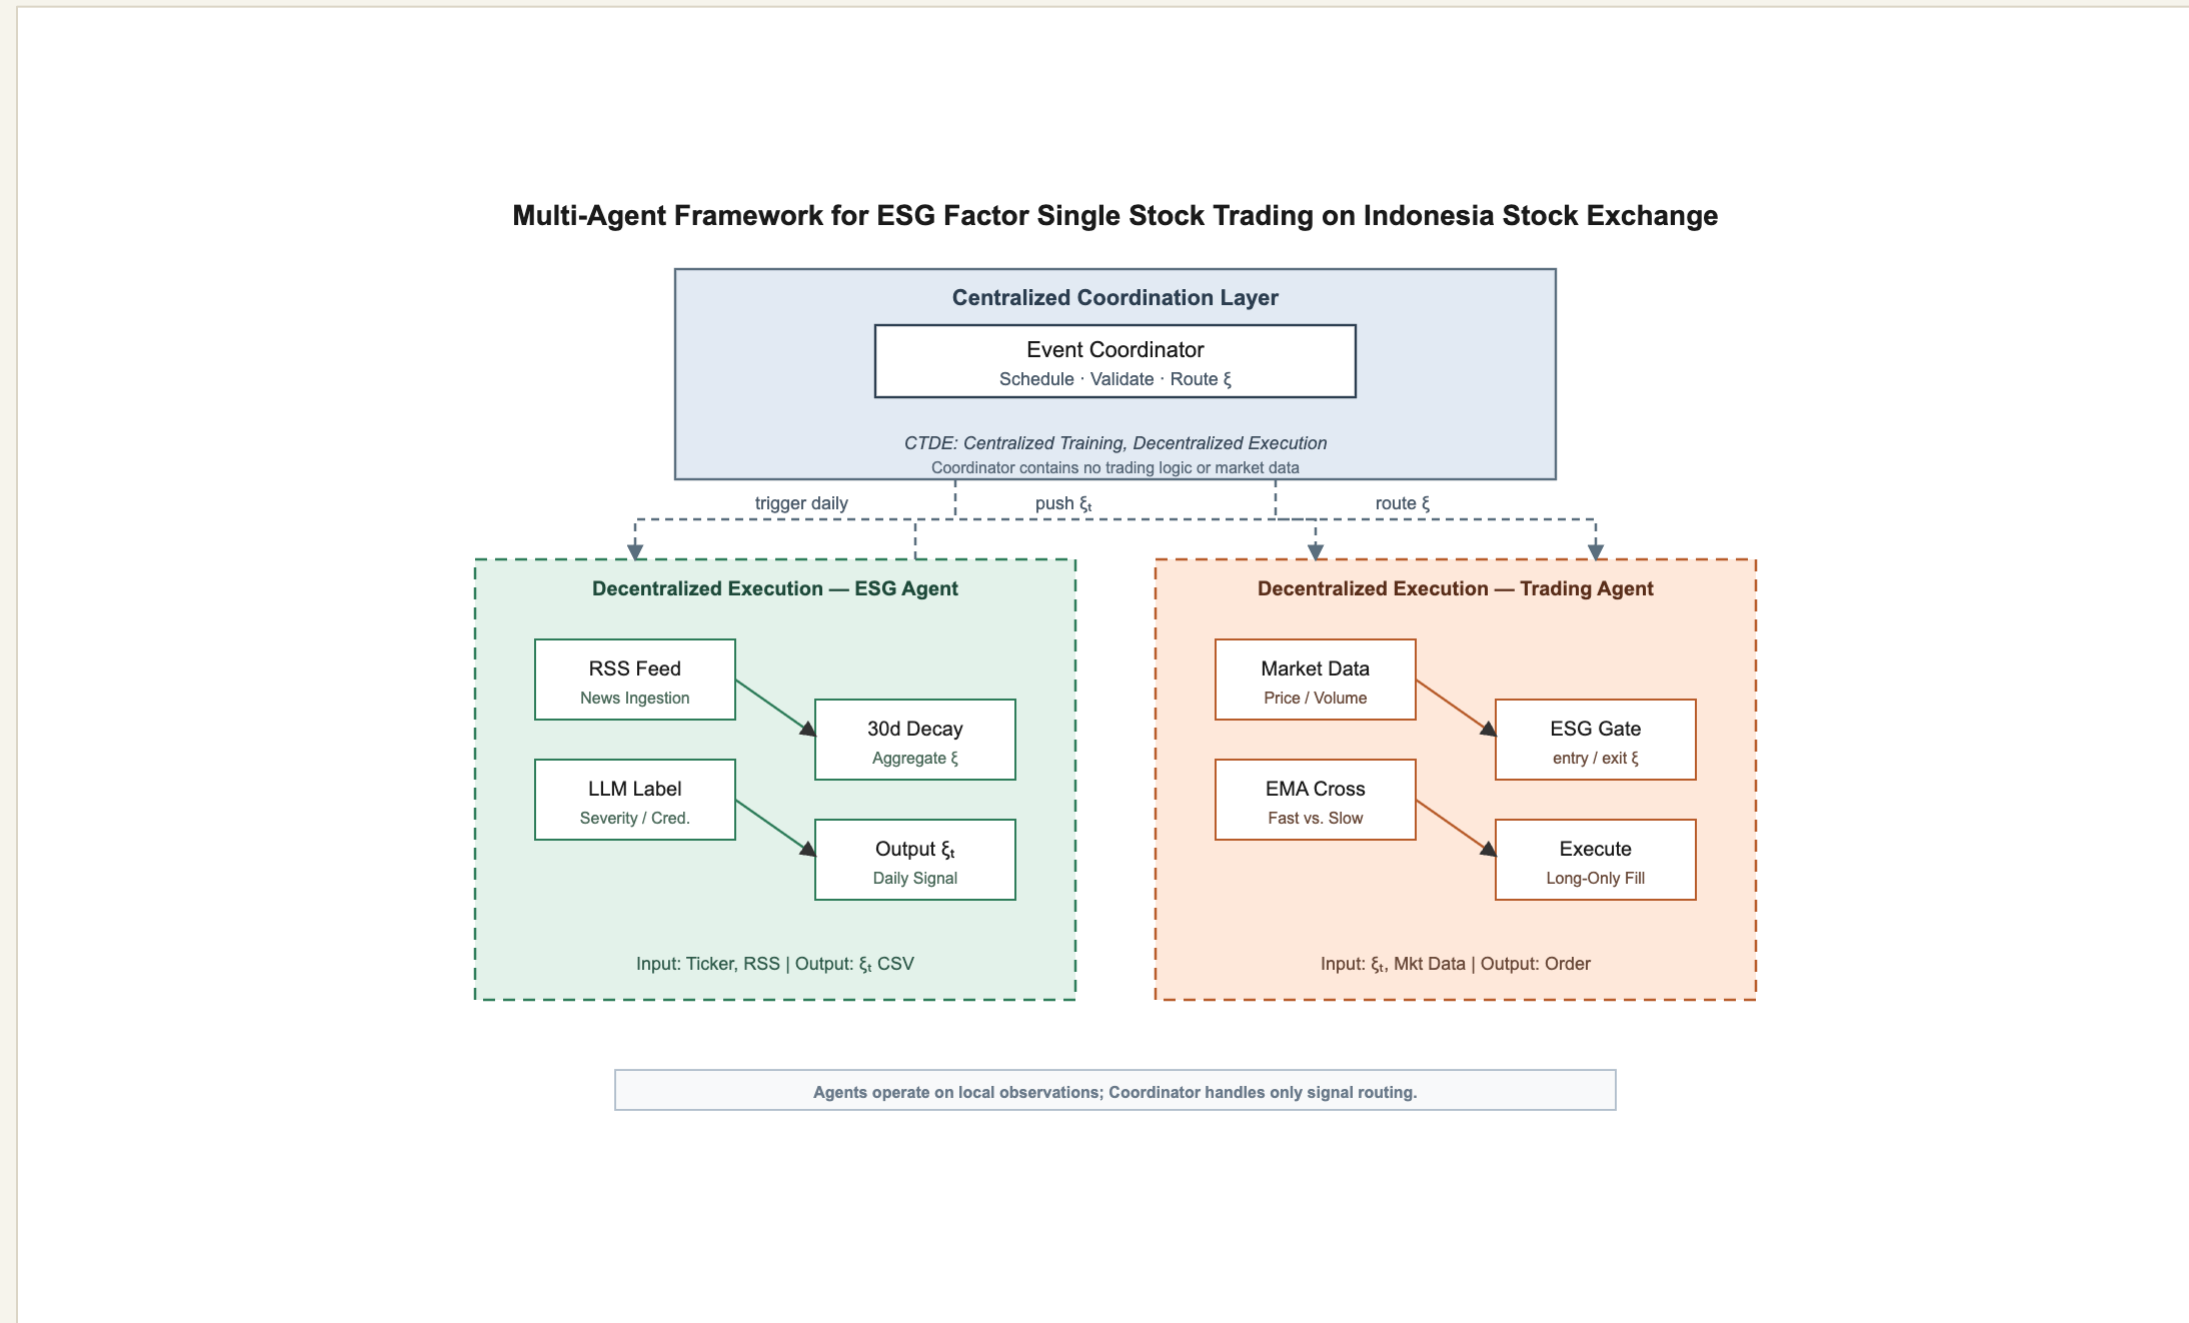
Computes EMA trend signals and applies ξ -based gating to entries and exits. **Never sees raw news.**

● EVENT COORDINATOR

Routing

Schedules the pipeline, validates signal integrity, enforces no-lookahead, and routes ξ by ticker. **No trading logic.**

Centralized coordination, decentralized execution



- **Coordinator**
Schedules, validates, routes ξ — no market data, no trading logic.
- **ESG Agent**
RSS → LLM label → 30-day decay → daily ξ_t .
- **Trading Agent**
Market data → EMA cross → ESG gate → long-only fill.

Agents act on local observations only.
Scales to multi-asset via ticker routing.

From news feed to a controversy signal

STEP 1 · FETCH

136

raw RSS items pulled for ANTM.JK across 2025.

STEP 2 · FILTER

20

source-backed candidates survive verification & dedup.

STEP 3 · LABEL

Severity Credibility Sentiment

three numeric fields per article from the LLM.

STEP 4 · AGGREGATE

ξ_t

30-day exponential decay into a daily intensity.

The controversy signal ξ , formalized

$$\text{weight}_{i,t} = \text{sev}_i \times \text{cred}_i \times \exp(-\ln 2 \cdot \text{age}_{i,t} / 30d)$$

$$\xi_t = 1 - \exp(-\sum_i \text{weight}_{i,t})$$

HALF-LIFE

30 calendar days — matched to the Indonesia ESG reporting convention.

RANGE

Bounded **0** → **1**; rises with recent severity & volume.

SENTIMENT ψ

Severity-weighted average computed too — held for future asymmetric gating.

EMA trend-following, gated by ESG

ENTRY BLOCK

Bullish EMA cross, but $\xi_t \geq \text{entry}_\xi \rightarrow$ entry ignored, no position opened.

FORCED EXIT

Position held, and $\xi_t \geq \text{exit}_\xi \rightarrow$ close at next open.

`constraint · entry_ξ < exit_ξ`

EXECUTION RULES

FILL

Next-open

COMMISSION

30 bps per side

SLIPPAGE

2 bps per side

CAPITAL

IDR 100,000,000

DIRECTION

Long-only · single position · no shorts

SECTION 03 · IMPLEMENTATION

The QBX Platform

Where the multi-agent framework becomes a declarative, validated, backtestable strategy graph.

The framework, as a validated strategy graph

Typed blocks

Market, indicator, external feature, comparator, logic, and action nodes — each a well-defined I/O contract, no opaque code.

ESG as a feature

The xi signal enters as a first-class `external_feature`, wired into the same gate — separation enforced structurally.

Audited execution

Long-only · next-bar-open fill · fee 30 bps · slippage 2 bps — every backtest runs this exact config.

QBX / Strategy Builder
Strategy Builder
 id_stock_esg_antm_temporal_signal-clone · ANTM.JK · 1h

Valid private · id_stock_esg_antm_tem... Edited LLM Save

Diagram Schema Yaml + Add Block Delete Edge Auto Layout 12 nodes

STRATEGY PICKER

ESG EVENT INPUT
 Where source-backed ESG news enters — the daily ξ signal feeding our AI-native research desk.

The diagram consists of the following nodes and connections:

- MARKET #10 (kline):** source: kline, outputs: open, high, low, close, volume.
- EXT #5 (esg):** external_feature, value.
- INDICATOR #12 (sma20):** ta.sma@1 · v1, period 20, outputs: sma.
- COMP #7 (esg_uncertainty_shock):** compare.gt@1 · v1, left: esg, right: 0.65.
- COMP #1 (close_above_sma):** compare.gt@1 · v1, left: sma, right: -0.2.
- COMP #2 (close_below_sma):** compare.lt@1 · v1, left: sma, right: 0.2.
- LOGIC #4 (entry_event):** logic.all@1 · v1, conditions: COMP #1, COMP #7.
- LOGIC #8 (exit_event):** logic.any@1 · v1, conditions: COMP #2, COMP #7.
- ENTER #3 (enter_long):** enter · long, trigger: LOGIC #4, filters: sizing, filter.all, filter.any.
- EXIT #9 (exit_long):** exit · long, trigger: LOGIC #8, filters: filter.all, filter.any.

Inspector

Strategy

- AME: Indonesia ANTM ESG Temporal Signal Clone
- icker: [empty]
- AAPL: AAPL market data for 1h.
- MEFRAME: 1h
- RECTION: long_only
- EXECUTION TYPE: long_only
- Management: 1
- Execution: configured
- SIGNAL TIMING: execution.bar_close_signal@1
- FILL MODEL: execution.next_bar_open@1
- INTRABAR PATH: -
- ORDER: execution.market_order@1
- FEE MODEL: execution.fee_bps@1
- FEE BPS: 30
- SLIPPAGE MODEL: execution.slippage_bps@1
- SLIPPAGE BPS: 2
- Execution is a strategy-level config used by backtests; it i...
- Selection: Select a node or edge

esg_sentiment_ok · gate entry esg_uncertainty_shock · force exit enter_long / exit_long · actions

fee 30 bps · slippage 2 bps · next_bar_open

SECTION 04

Experiments & Results

A 3,825-run parameter grid, seven ablation conditions,
and a temporal placebo.

A full grid sweep and seven ablations

PARAMETER GRID

FAST EMA 5 · 8 · 10 · 12 · 16 · 20 · 30 · 40 · 50

SLOW EMA 30 → 200 (fast < slow)

POSITION FRAC 0.25 · 0.50 · 0.75

ENTRY ξ 0.2 · 0.35 · 0.5 · 0.65 · 0.8

EXIT ξ 0.55 · 0.7 · 0.85 · 0.95

3,825 VALID COMBINATIONS

SEVEN ABLATION CONDITIONS

price_only EMA trend only — ESG off

entry_only block entries, no forced exit

exit_only forced exits, no entry block

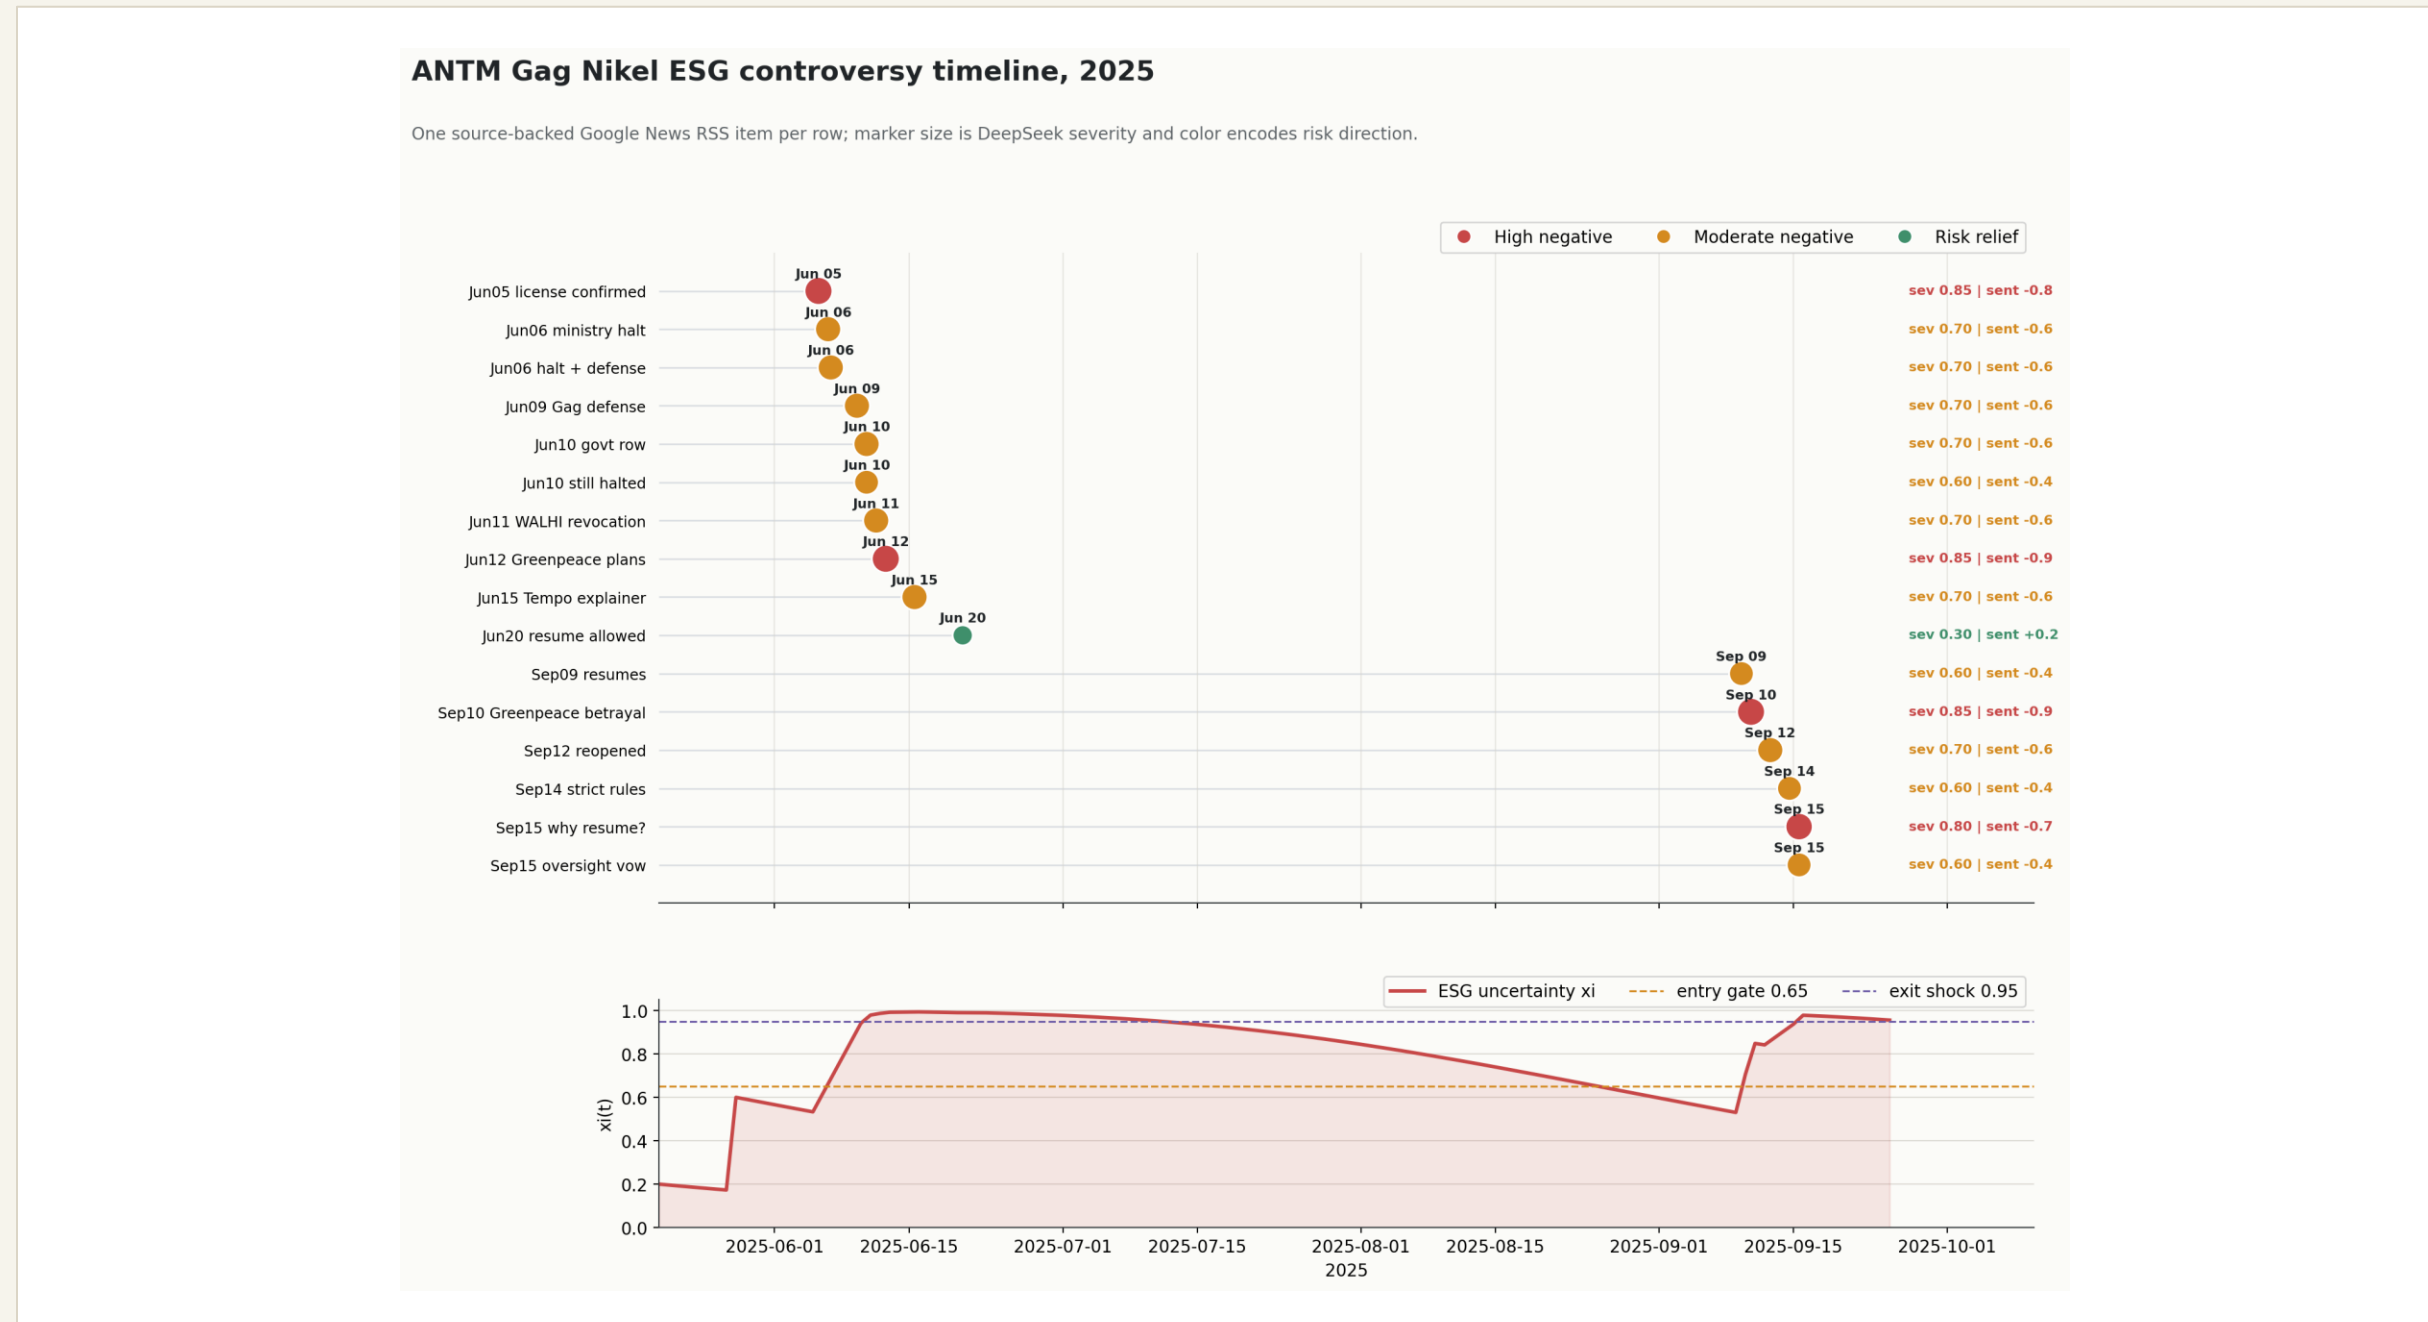
full_actual both gates active

full_lag5 ξ delayed 5 bars

full_shuffled temporal placebo — order randomized

full_inverted diagnostic — signal inverted

The Gag Nickel controversy, as a signal



Two dense negative clusters — **June** and **September** — bracket a single risk-relief event. Marker size is LLM severity; colour is risk direction.

The resulting ξ_t spikes, plateaus, and decays — crossing the **entry gate** and **exit shock** thresholds.

When ξ is elevated, the strategy stays out of the market.

ESG gating roughly doubles the Sharpe

CONDITION	MEAN SHARPE	BEST SHARPE	% POSITIVE	MEAN FILLS	MEDIAN MAX DD
full_actual	1.064	3.339	83.1%	2.75	-0.035
full_lag5	1.079	3.308	96.6%	2.59	-0.038
price_only	0.539	1.765	70.7%	3.93	-0.122
entry_only_actual	0.459	2.040	79.1%	2.31	-0.073
exit_only_actual	0.391	2.095	59.6%	53.33	-0.126
full_inverted	-0.016	1.845	31.5%	3.40	-0.111
full_shuffled	-0.593	1.852	33.4%	23.93	-0.120

Isolating the ESG contribution

COMPARISON (FULL_ACTUAL - ...)	Δ MEAN SHARPE	Δ MEDIAN SHARPE	FULL ACTUAL BETTER	Δ MEAN FILLS
- price_only	+0.526	+0.536	73.0%	-1.19
- full_shuffled	+1.657	+1.657	95.7%	-21.19
- full_inverted	+1.080	+1.116	79.2%	-0.66
- exit_only_actual	+0.673	+0.800	79.2%	-50.58
- entry_only_actual	+0.605	+0.775	59.1%	+0.44
- full_lag5	-0.015	0.000	38.6%	+0.15

A representative parameterisation

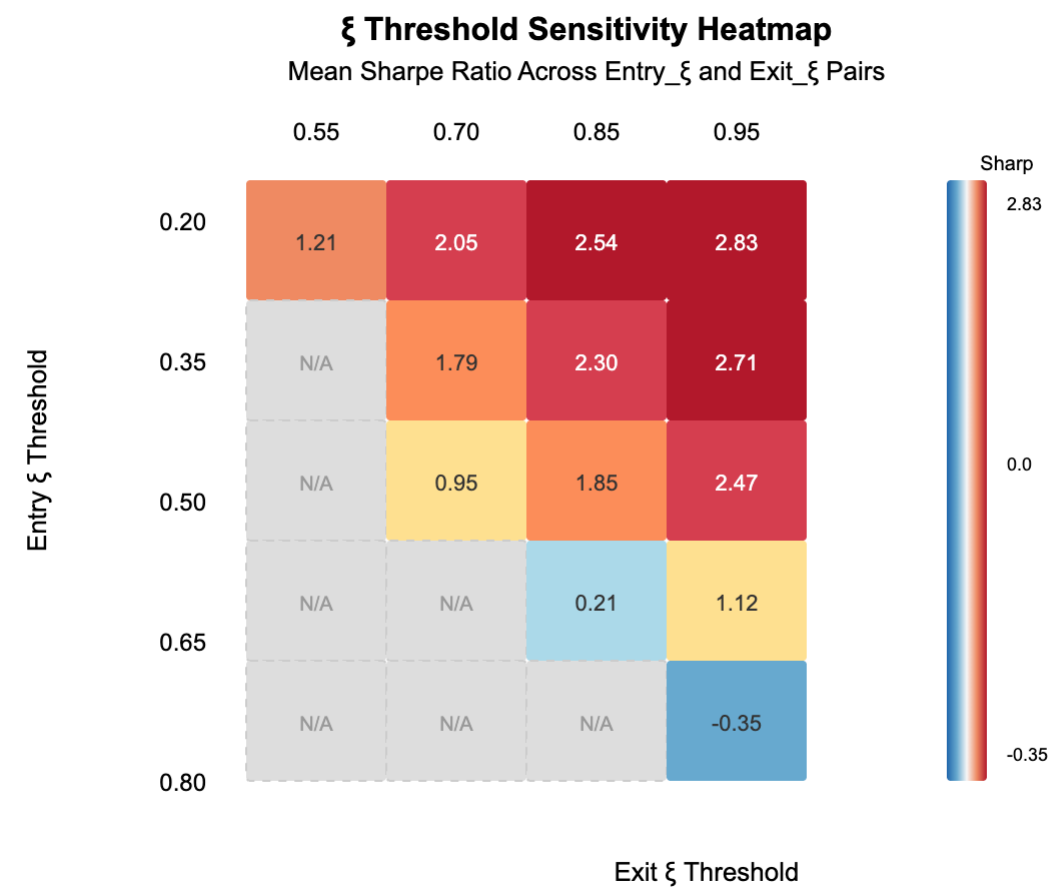
CONDITION	RETURN	SHARPE	MAX DD	FILLS	EXPOSURE
full_actual	+111.4%	3.339	-0.071	4	0.216
full_lag5	+96.2%	2.990	-0.093	4	0.237
price_only	+63.6%	1.639	-0.184	2	0.598
exit_only_actual	+68.4%	1.903	-0.160	34	0.530
full_shuffled	+5.1%	0.330	-0.304	34	0.449

FULL_ACTUAL

3.34 net Sharpe

4 fills **-7.1%** max DD **22%** exposure

Robust across thresholds — but not uniform



The ESG signal alters decisions across a **wide region** of (entry_ξ, exit_ξ) pairs — most of the surface is favourable.

But entry thresholds that are **too strict** (e.g. 0.8) suppress profitable late-year momentum entries, because ξ stays elevated after the controversy clusters.

A risk conditioner that needs calibration — not a magic switch.

SECTION 05

Discussion & Conclusion

What the evidence says about ESG, market efficiency,
and auditable financial AI.

What the evidence shows

A ESG is a material risk conditioner

A +0.53 Sharpe premium — economically significant and parametrically robust.

C The architecture scales

A minimal Coordinator here becomes the routing hub for a multi-asset resource portfolio.

B CTDE enables clean attribution

Decoupled agents make ablation a routing toggle; both gates cut fills from 53 to under 3.

D LLM labelling is viable & auditable

Source-backed signals; the shuffle test proves value lies in timing, not distribution.

Frictional, semantic weak-form efficiency

Weak-form holds

The price-only baseline fails — the IDX efficiently impounds simple quantitative signals.

Semi-strong refuted

Excess returns from public ESG news survive 60 bps costs — and vanish under a temporal shuffle.

A structural window

High-context narratives diffuse over weeks; staying out of the market materially improves net returns.

Reading the news, trading with discipline — and finding the market's blind spot.

A modular, auditable multi-agent framework turns unstructured ESG news into actionable risk-gating — and beats a crowded retail market even after 60 bps costs.

Dingyue Kan · Yichong Bai — HK PolyU

Jerry Cao — Hang Seng University of Hong Kong

Thank you Key refs · Fama 1970 · Hong & Stein 1999 · Lowe 2017 · Vadori 2024

Computation for this work is backed by **UCAIC**, The Hong Kong Polytechnic University.