AI Sentiment Market Prediction for Crypto Trading with Proof-of-Signal Verification

A hybrid approach combining technical indicators and sentiment analysis to generate verifiable, tamper-evident trading signals for volatile cryptocurrency markets.

Andrei Martynenko K1946577

| Faculty of Engineering, Computing and Environment | Kingston University

VERIFIABLE AI-POWERED BLOCKCHAIN-READY

Research Foundation and Objectives

The Challenge

Cryptocurrency markets are highly volatile, and AI trading platforms often operate as "black boxes" without verifiable records. This lack of transparency undermines trust and reproducibility.

Core Research Questions

H1 (Auditability): Can a trading system provide tamper-evident records of signal generation?

H2 (Signal Quality): Does the hybrid approach accurately predict crypto market direction?

Primary Aim

Develop a crypto signal system producing interpretable BUY/SELL/HOLD decisions with tamper-evident proof for enhanced trust and auditability.

Stakeholder Benefits

- Traders gain transparent evidence
- Researchers get reproducible evaluation
- Fintech adopts proof mechanisms

Research Hypotheses

1

Hybrid Signal Superiority

Combining sentiment analysis with technical indicators produces higher-quality trading signals compared to single-source approaches (sentiment-only or technical-only).

Focus: Prediction accuracy and actionable signal generation.

2

Proof-of-Signal Integrity

Implementing deterministic proof hashes with optional blockchain anchoring improves signal verifiability, integrity, and audit capabilities.

Focus: Reproducibility, tamper-evidence, and accountability.

These hypotheses address both *predictive performance* and *system trustworthiness*—essential for institutional adoption.

System Architecture and Design Methodology

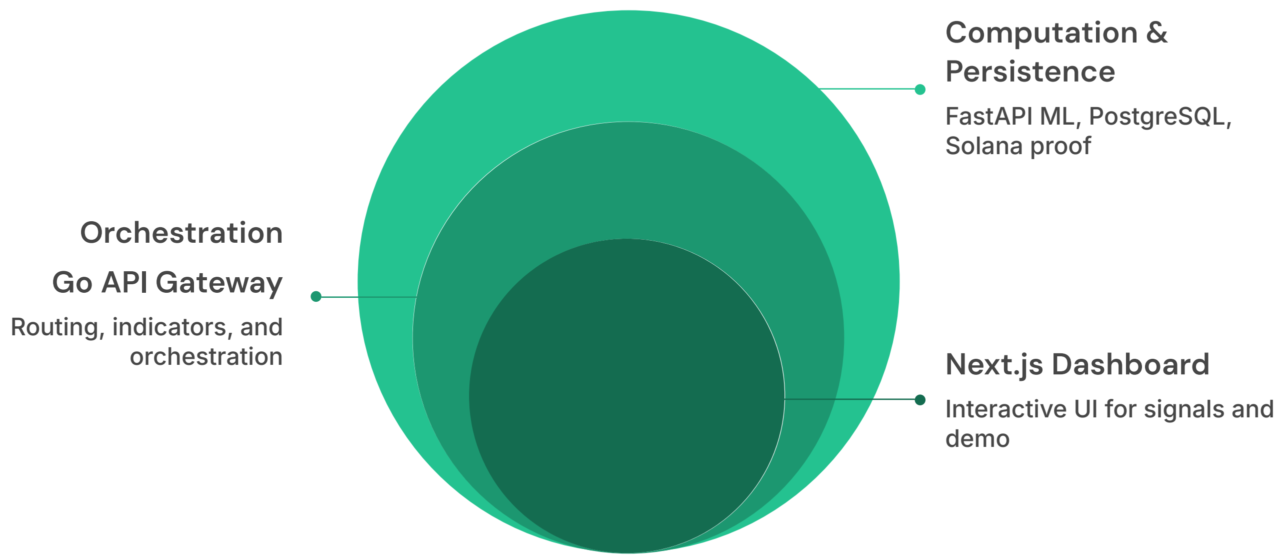

Architectural Rationale

The system uses a **microservice-style**, layered architecture separating presentation (Next.js), orchestration (Go gateway), and computation (FastAPI ML service). This improves maintainability, testability, and allows each component to evolve independently. Development followed an iterative vertical-slice approach to achieve early end-to-end integration and reduce interface risk.

Next.js Dashboard

User-facing interface that visualizes signals and proof metadata (authentication via Supabase when enabled).

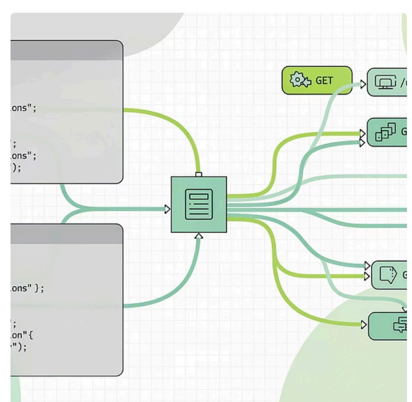

Go API Gateway

Provides a stable REST API for the frontend, routes requests to downstream services, and supports optional background processing (institutional runner) and DB-backed endpoints.

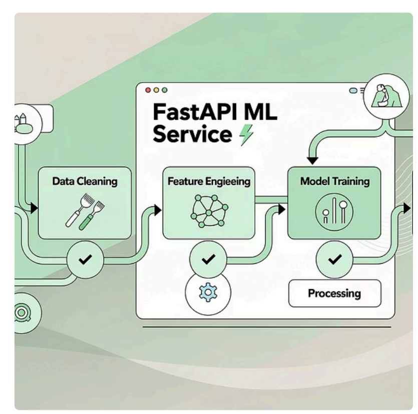

FastAPI ML Service

Runs sentiment analysis (FinBERT-style), hybrid scoring, and proof-of-signal generation via deterministic hashing (optionally anchored on Solana).

Data Pipeline and Feature Engineering

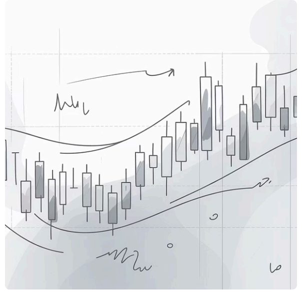

Market Data Inputs

OHLCV (Open, High, Low, Close, Volume) time-series data capturing price movements and trading activity across multiple timeframes.

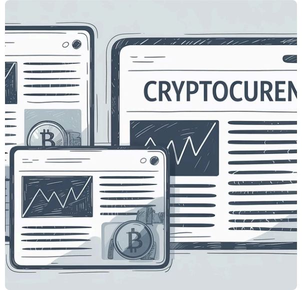

News Text Sources

Crypto-related news articles and social media content providing real-time sentiment signals and market narrative context.

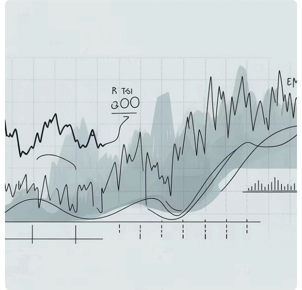

Technical Features

Classic indicators including Exponential Moving Averages (EMA), Relative Strength Index (RSI), and MACD for trend identification.

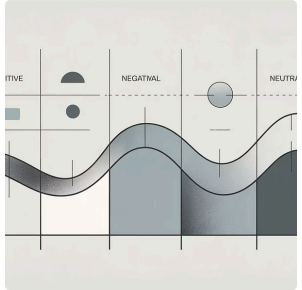

Sentiment Features

FinBERT-powered classification producing positive/negative/neutral labels with confidence scores, augmenting technical analysis.

Output: Hybrid trading action (BUY/SELL/HOLD) with confidence scores and human-readable reasoning for each signal.

Proof-of-Signal: The Key Innovation

Verifiable Integrity

Every signal includes a deterministic SHA-256 proof hash, making predictions tamper-evident and reproducible.

Core principle: Identical input payloads *must* produce identical proof hashes, enabling independent verification.

01

Canonical JSON Payload

Signal metadata (timestamp, asset, action, confidence, features) formatted consistently.

02

SHA-256 Proof Hash

Cryptographic fingerprint computed from payload, stored with every signal output.

03

Optional Solana Anchoring

Proof hash submitted to Solana devnet, returning transaction signature as external timestamp reference.

04

Audit Trail

Complete traceability: reconstruct signal conditions and verify authenticity at any future point.

Implementation: See `ml_service/solana_layer.py` functions `hash_signal()` and `send_proof()` for technical details.

Interpretation: This mechanism supports **H2** by making each generated signal independently verifiable. Because the proof is a deterministic SHA-256 hash of a canonical signal payload, any party can recompute the hash and confirm the signal has not been altered. Any change to the payload produces a different hash, making tampering detectable. When enabled, a Solana devnet transaction signature provides an external timestamp reference that strengthens auditability.

Evaluation Methodology and Results

Offline Backtest Design

Signal quality was evaluated using historical Binance OHLCV data for BTCUSDT and ETHUSDT across October 2025 to January 2026 (1-hour candles). Ground-truth labels were derived from forward returns: for each timestamp t, the return over horizon $H = 24$ hours was computed, with labels assigned as BUY if return ≥ +2%, SELL if return ≤ −2%, and HOLD otherwise. This explicit labelling rule provides reproducible ground truth whilst filtering noise through the threshold parameter.

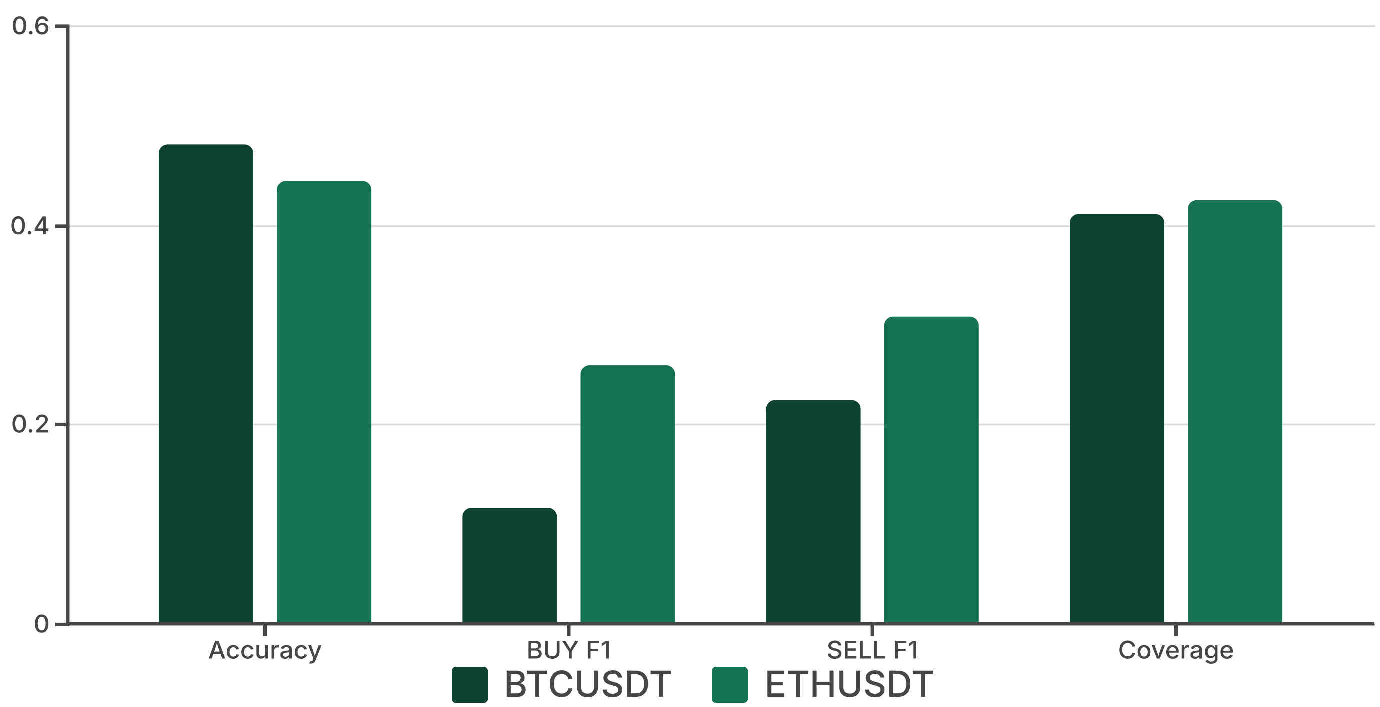

Performance Summary

Across 2,266 hourly signals per asset, BTCUSDT achieved 48.1% accuracy with stronger SELL detection (F1 = 0.225) than BUY (F1 = 0.117). ETHUSDT showed 44.4% accuracy with more balanced class performance (BUY F1 = 0.259, SELL F1 = 0.309).

Coverage metrics indicate the system emits actionable BUY or SELL signals in approximately 41–43% of periods, avoiding excessive HOLD classification whilst maintaining directional discipline.

Interpretation: The backtest provides evidence for **H1** because the system shows **measurable predictive structure** when classifying future 24-hour market moves into BUY/SELL/HOLD under a reproducible labeling rule, and performance is evaluated with **precision/recall/F1** (not just accuracy), revealing realistic differences between BTC and ETH. However, these results **do not prove profitability** (fees, slippage, sizing, and risk controls are not modeled)

Results and Achievements

End-to-End Prototype

Fully functional prototype demonstrating the pipeline across sentiment analysis, technical indicator scoring, and hybrid signal generation via API, with a web dashboard for visualization.

Proof Hash Integration

Signals support a tamper-evident SHA-256 proof hash, and optional Solana devnet anchoring can be enabled to return a transaction signature as an external timestamp reference.

Reproducibility Demonstrated

Deterministic proof generation validated: identical signal payloads consistently produce identical proof hashes, enabling independent verification.

Important Context: Proof-of-Concept Scope

This system demonstrates technical feasibility and validates the core hypotheses, but it is not production-ready for live trading. Key constraints include API rate limits, data quality variability, and model limitations under rapidly changing market regimes; overfitting risk and real execution factors (fees, slippage, latency) are not fully addressed in this prototype phase.

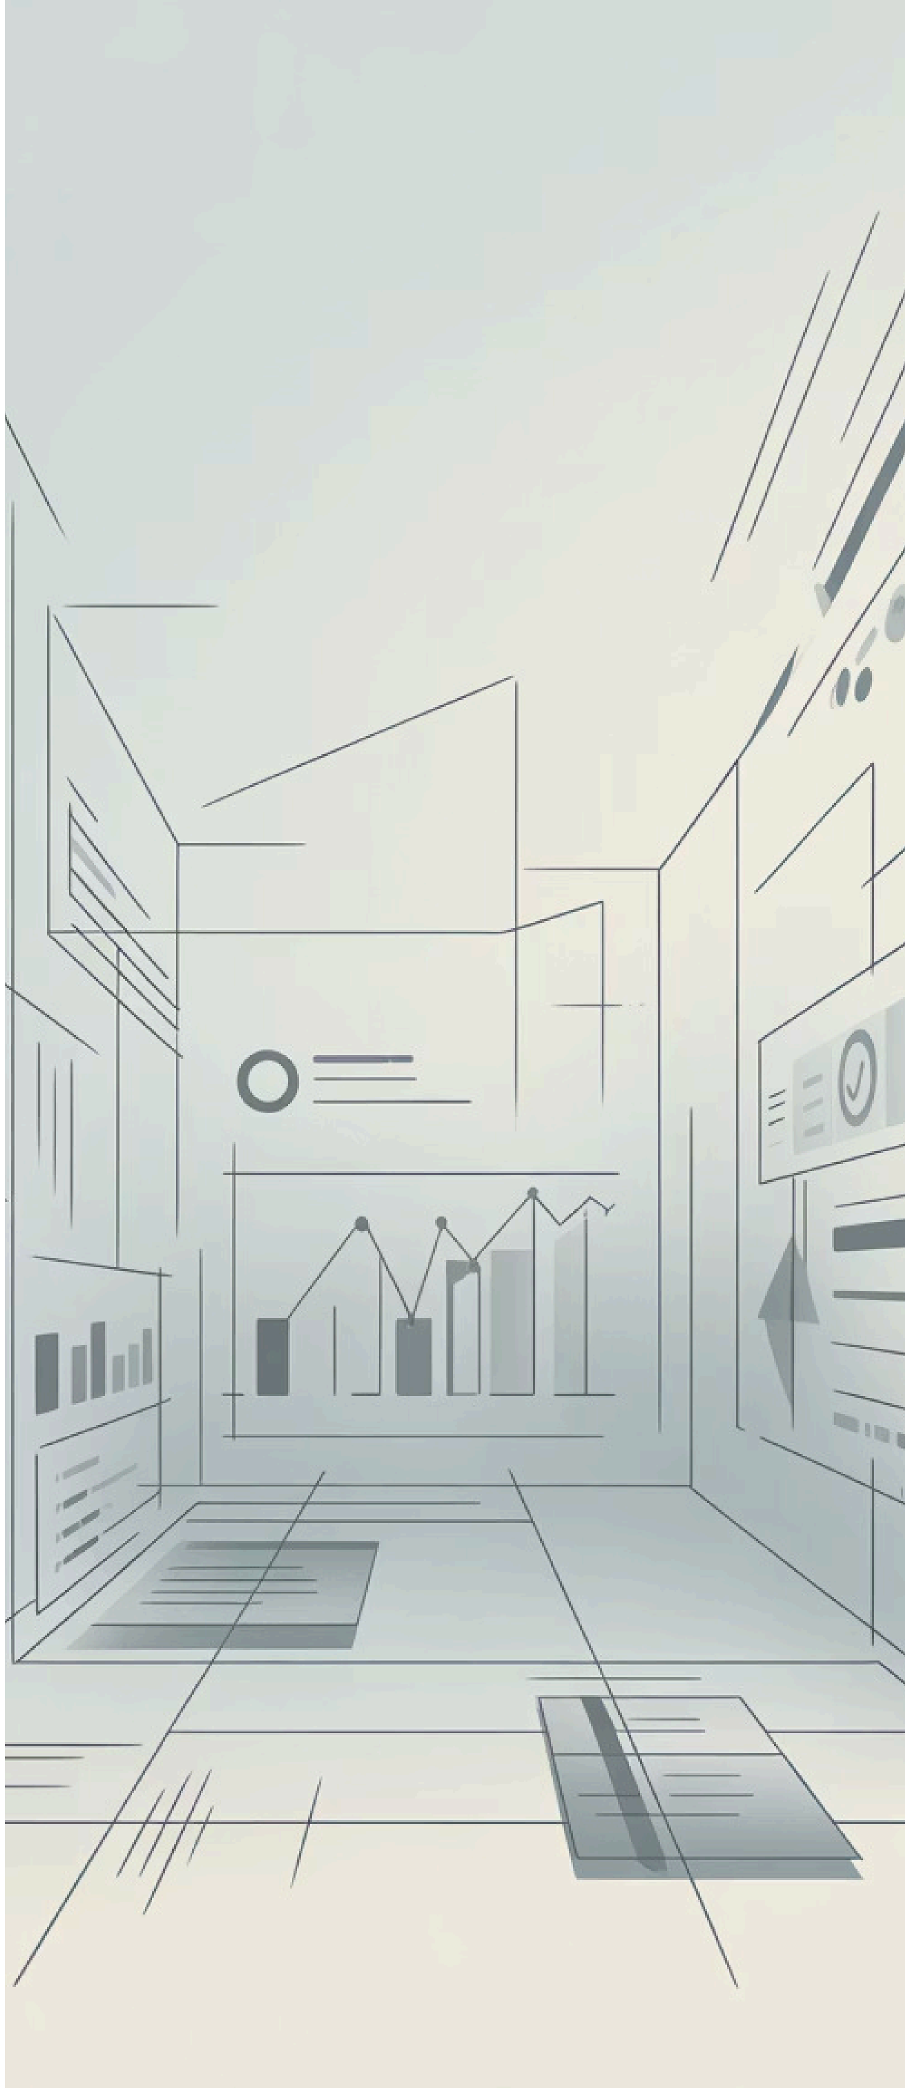

Live Demonstration

Five-Step Demo Flow

1

Open Dashboard

Navigate to signals and news monitoring interface.

2

Trigger Analysis

Fetch fresh news and technical indicators for BTC/ETH.

3

View Signal Output

Display returned signal: action, confidence, reasoning.

4

Inspect Proof Hash

Show deterministic SHA-256 hash in API response and UI.

5

Verify Blockchain Anchor

Present Solana transaction signature.

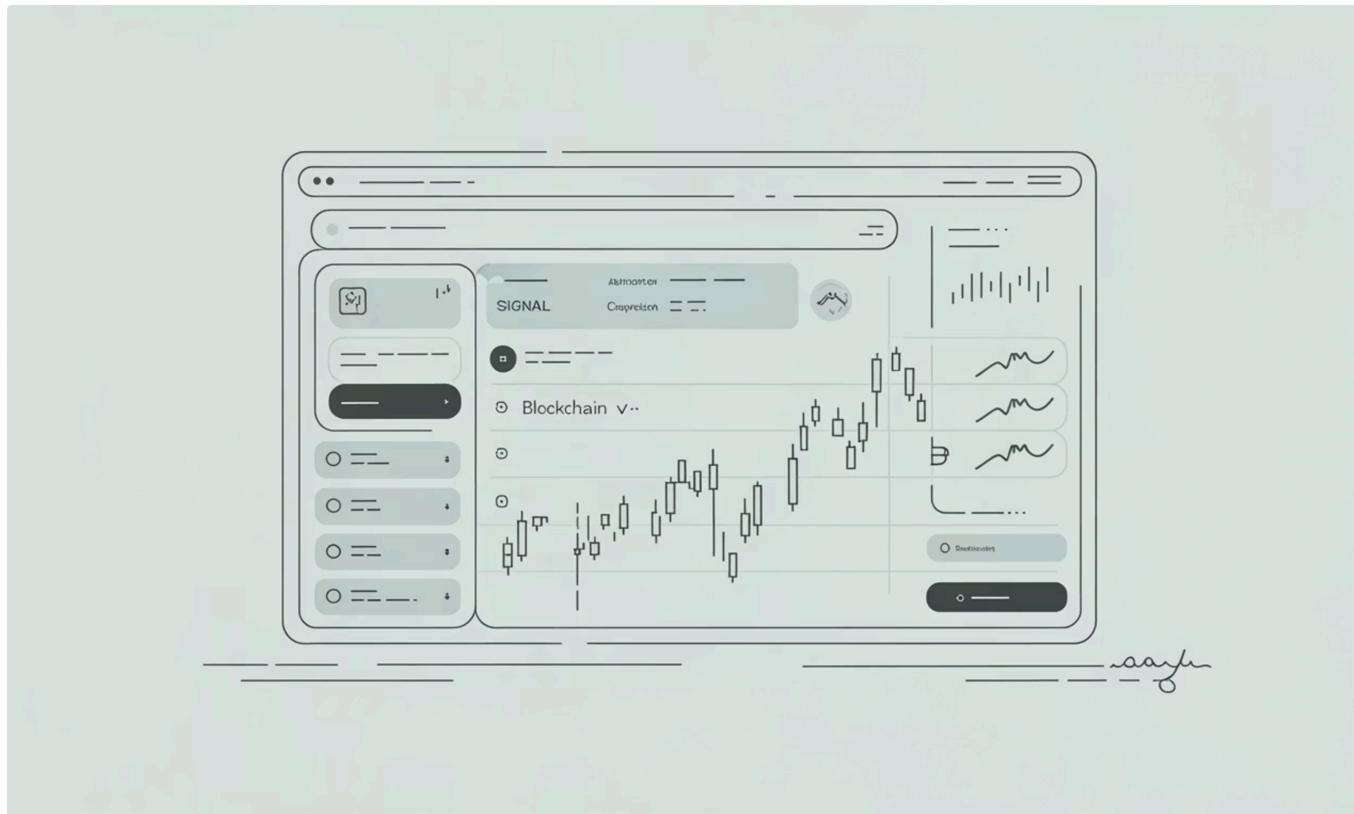

> "This demonstrates the complete pipeline and proves signal verifiability: the proof hash uniquely represents this signal payload, enabling independent audit."

Critical Analysis and Future Directions

Current Limitations

- **Sentiment noise:** Sarcasm detection, source credibility assessment, and context disambiguation remain challenging.
- **Non-stationary markets:** Regime shifts and structural breaks can invalidate trained model assumptions rapidly.
- **Backtest vs reality gap:** Slippage, transaction fees, latency, and liquidity constraints not captured in simulation.

Future Research Directions

- **Baseline comparisons:** Systematic evaluation against sentiment-only, technical-only, and naive strategies.
- **Walk-forward validation:** Time-series cross-validation to better estimate out-of-sample performance.
- **Execution simulator:** Realistic transaction cost modelling with risk management and position sizing.
- **Enhanced provenance:** Cryptographically signed proofs, richer metadata logging, and database-first audit trails.

Key insight: Prediction accuracy differs fundamentally from profitable trading. Transaction costs, market impact, and risk management transform signal quality into actual returns.