

A system for organizing uncertainty.

Crene is a **structured uncertainty system** for institutional data buyers. It decomposes the macro, AI, and fiscal questions that have no clean market price into falsifiable components, scored daily across a four-model frontier ensemble, with calibration disclosed openly. **The data is the moat.**

API + DATASET

FOUR-MODEL ENSEMBLE

CALIBRATION DISCLOSED

OPEN-SOURCE FRAMEWORK

BOOTSTRAPPED

THE PROBLEM

The questions that move portfolios have no price.

Funds must position around regime shifts, AI labor displacement, fiscal stress, and demographic change. Markets price only what is liquid and near-dated. Everything structural is left to narrative: analyst notes, conviction, vibes. No probability, no calibration, no falsifiable structure on the questions that actually decide returns.

THE PRODUCT

A dataset funds can license. A probability they can audit.

Crene decomposes major unresolved questions into falsifiable components, scored daily across a four-model ensemble and calibrated against resolved outcomes. The buyer-facing objects stay simple: Scenario, Cluster, Factor, with Event as the calibration atom. Delivered by API and dataset license.

4 scenarios 464 components 600 pathways 3 clusters 350 factors

THE WEDGE

01 / NOW

Dataset

Structured uncertainty data, listed for distribution through institutional channels. Funds buy differentiated, calibrated signal they cannot build in-house.

02 / NEXT

API + Framework

Probability endpoints for live and resolved events. signal-tracker open-sourced on PyPI: the framework is free, the data is the moat.

03 / END STATE

Calibration layer

The reference probability source for the strategic questions markets cannot price. The place institutions check before outcomes resolve.

CALIBRATION · DISCLOSED
OPENLYLISTED FOR DISTRIBUTION · NEUDATA / EAGLE ALPHA / MONDA · PRE-
CONTRACT

0.114

EXTERNAL ENSEMBLE BRIER
(N=811, MARKET-VALIDATED)

0.2325

MACRO-ONLY BRIER
(N=242)

1,100+

RESOLVED EVENTS
TRACKED

4

FRONTIER MODELS
SCORED DAILY

Honest tiering. The 0.114 validates the ensemble against market-priced questions; it is not product proof. The macro product record (0.2325, modest skill over base rate) and the scenario layer are still accruing, because those resolve years out. Calibration is published, not claimed.

WHY NOW

AI can finally decompose and score uncertainty at scale, every day, across an ensemble. The questions that move portfolios are exactly the ones markets cannot price. Whoever organizes that, and proves it with calibration, owns the layer.

BUILDER VELOCITY · RECENT

- Four-model daily ensemble live: Claude Haiku, GPT-4o mini, Gemini 2.5 Flash Lite, Grok 4 Fast.
- Institutional package shipped: Product Sheet, Datasheet, Methodology, Data Guide.
- signal-tracker open-sourced to PyPI (MIT).
- Infra consolidated: crawler, content backend, and frontend on a single deployment.
- Honest calibration tiering published to the live product.

FOUNDER **Stephen Lee, Founder & CEO.** Seven years in institutional finance at Goldman Sachs and Credit Agricole, now self-taught full-stack, shipping daily across crawler, backend, and frontend. Based in San Francisco, CA, USA and Seoul, South Korea.

WHAT'S AVAILABLE

API probability endpoints

Dataset licensing

Methodology + Data Guide

Pilot conversations open