



AUTONOMOUS-FOX
LABORATORIES





AUTONOMOUS-FOX
LABORATORIES

Foundational AI for Financial Markets

Who We Are



- UK-based deep-technology research lab
- Conducting foundational artificial intelligence research in behavioural modelling for financial markets
- Building infrastructure-level models of market behaviour
- 12+ years of academic and applied machine learning in trading and market making
- Research-led | Pre-revenue | IP-focused

A Fundamental Gap In Market Modelling

Deep learning produced reusable representation layers that transformed language, vision, and biological science.

No equivalent shift has occurred in finance.

Production systems still model markets through narrow, task-specific approaches focused on pattern recognition in raw observable data.

What is missing is a coherent, continuously updating representation of how assets are behaving, individually and in relation to one another, as conditions evolve.

Why Trading Systems Fail to Scale

The Reality of Markets

Financial markets are among the most difficult real-world systems for machine learning:

- Non-stationary
- Adversarial
- Extremely low signal-to-noise ratios
- Complex, time-varying dependencies

The Consequence

Modern trading systems evolved as collections of narrow, task-specific models:

- Each learns its own limited view of the market
- No transfer between assets or functions
- Complexity grows faster than performance
- Constant retraining and manual intervention

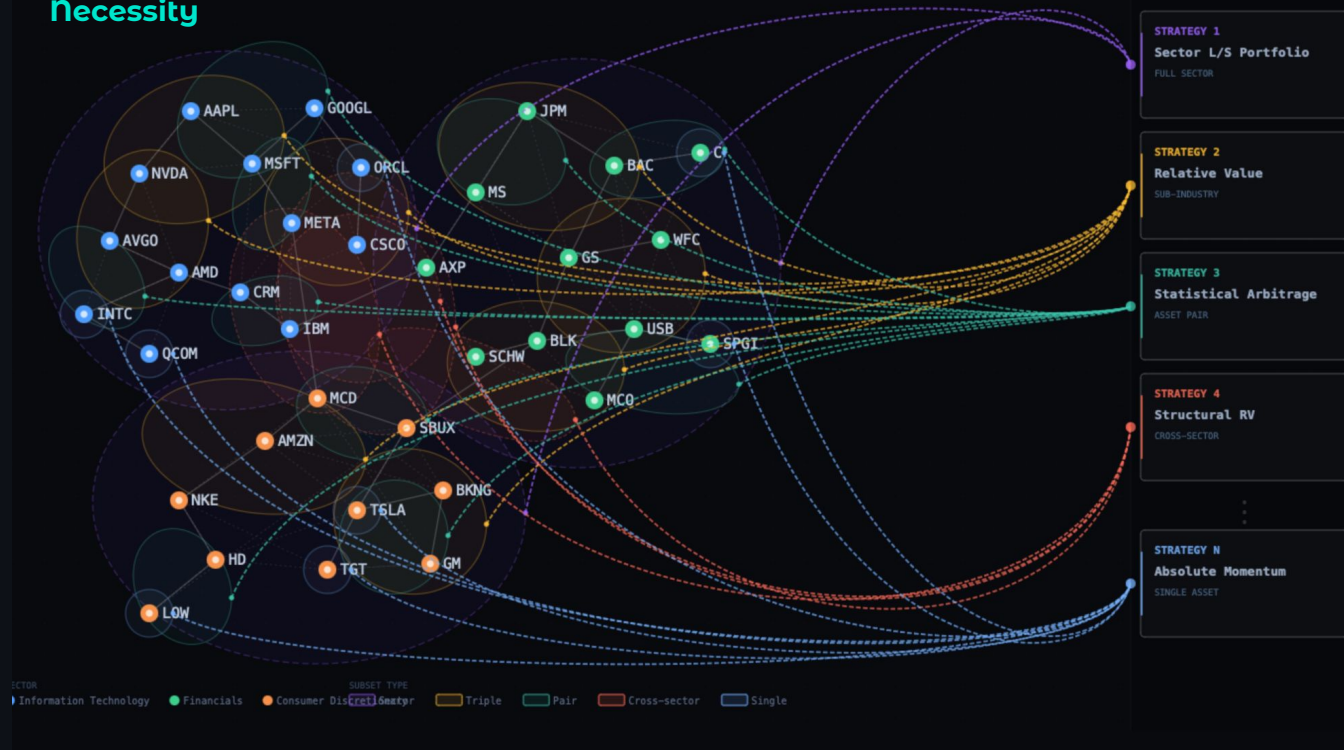
The Scaling Bottleneck

- No shared, evolving behavioural representation

Without a shared representation layer, trading systems do not scale.

The Fragmentation Problem

Illustrative US Equities Trading Stack - Fragmented by Necessity



Many models. No shared context.

- Each strategy sees only its own bounded context.
- Relationships outside the model's context are invisible.
- No transfer between strategies or asset classes.
- Complexity and redundancy grows faster than scale and performance..
- When market conditions shift, every model must be retrained because its context has changed

Our Research Philosophy

That the behaviour of individual assets can be decomposed into a vocabulary of discrete behavioural building blocks.

Each building block captures a distinct mode of market behaviour. That structure is not imposed. It is learned directly from the data.

LANGUAGE

Continuous speech → discrete subword tokens

The abstraction that made foundation models possible at scale.

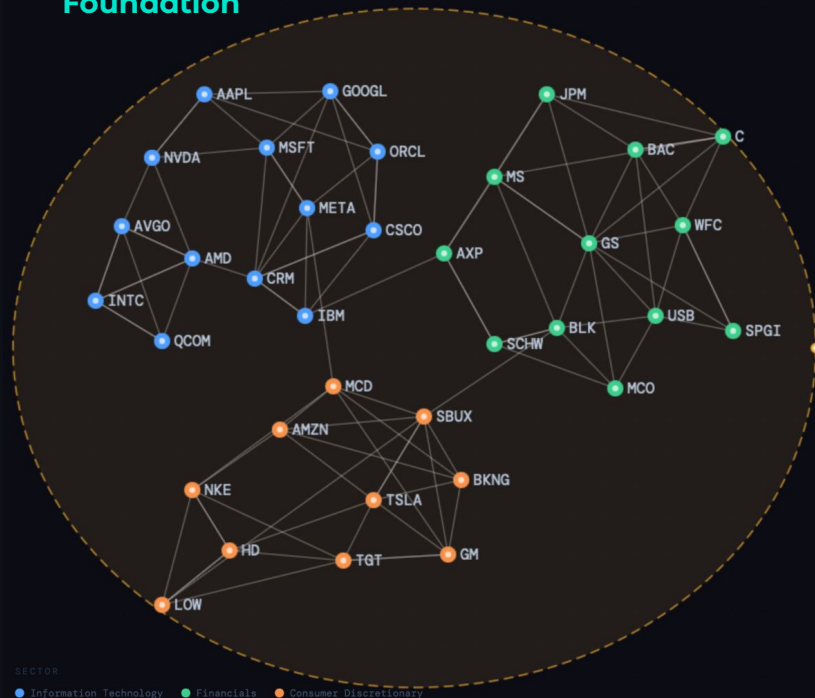
Continuous molecular structures → discrete structural representations

The abstraction that enabled breakthroughs in protein structure prediction.

Get the abstraction right, and the rest of the modelling stack can scale on top of it.

What We Are Building

Illustrative US Equities Trading Stack · One Shared Foundation



BEHAVIOURAL MODELLING
SUBSTRATE

One model.
Full market context.

A single, task-agnostic
representation of how assets
behave – individually and in
relation to one another –
continuously updated as
conditions evolve.

STRATEGY 1
Sector L/S Portfolio
FULL SECTOR

STRATEGY 2
Relative Value
SUB-INDUSTRY

STRATEGY 3
Statistical Arbitrage
ASSET PAIR

STRATEGY 4
Structural RV
CROSS-SECTOR

STRATEGY N
Absolute Momentum
SINGLE ASSET

One foundation. Full market context:

- One substrate learns the full market — every asset, every relationship.
- Continuously updated as conditions evolve.
- Every strategy draws from the same representation.
- Downstream models specialise representations, not raw data.
- Shifting market conditions tracked centrally — retraining overhead minimised.

The Deep Learning Problem in Finance



Deep learning transformed language, vision, and biological science. It has not yet had an equivalent impact in finance.

The Scale Approach

More data, more parameters, more compute, more capacity

In finance, more capacity gives you more rope to memorise noise. Deep networks overfit to patterns that were never real, and collapse when conditions change.

The domain violates core assumptions of standard architectures. Signal is weaker. Non-stationarity and adversarial feedback are intrinsic. Existing methods were not designed for this.

Our Thesis

Deep learning's value in finance does not come from scale, depth, or capacity.

It comes from the architectural flexibility that it gives us — the flexibility to control how models learn structure in noisy, continually changing environments.

Recent advances in deep learning now provide the architectural tools to make this possible.

In finance, structure beats scale.

What Has Changed

Three established areas of deep learning, now mature enough to give us sufficient and flexible control over the learning of structure in non-stationary, low-signal environments. A decade of proprietary behavioural modelling research provided the domain foundation to recognise when these methods were ready — and to make them work.

REPRESENTATIONAL STRUCTURE

Self-Supervised Representation Learning

Controlling what a model learns from raw data

The principle behind foundation models in language and vision. Models discover compact, reusable structure from data without requiring labelled examples.

TEMPORAL STRUCTURE

Sequence Modelling

Controlling how a model captures dynamics across time

The basis for modelling non-stationary physical and financial systems. Architectures that learn temporal dependencies and how patterns evolve, persist, and transition.

RELATIONAL STRUCTURE

Geometric Deep Learning

Controlling how a model organises dependencies between entities

Behind state-of-the-art weather forecasting models such as GraphCast and GenCast. Methods for learning over relational structures such as graphs and networks.

The Research Architecture

How we shape the learning of structure in non-stationary, low-signal environments

1 Discrete Tokenisation

Self-Supervised Representation Learning

Controlling what the model learns from raw data. Continuous market behaviour is compressed into a discrete vocabulary of learned representations, giving us direct control over what structure is preserved and what noise is discarded.

2 Fine-Grained State Inference

Sequence Modelling

Controlling how the model captures dynamics over time. Behavioural states and their temporal transitions are learned rather than specified, discovering structure at a resolution conventional approaches cannot achieve.

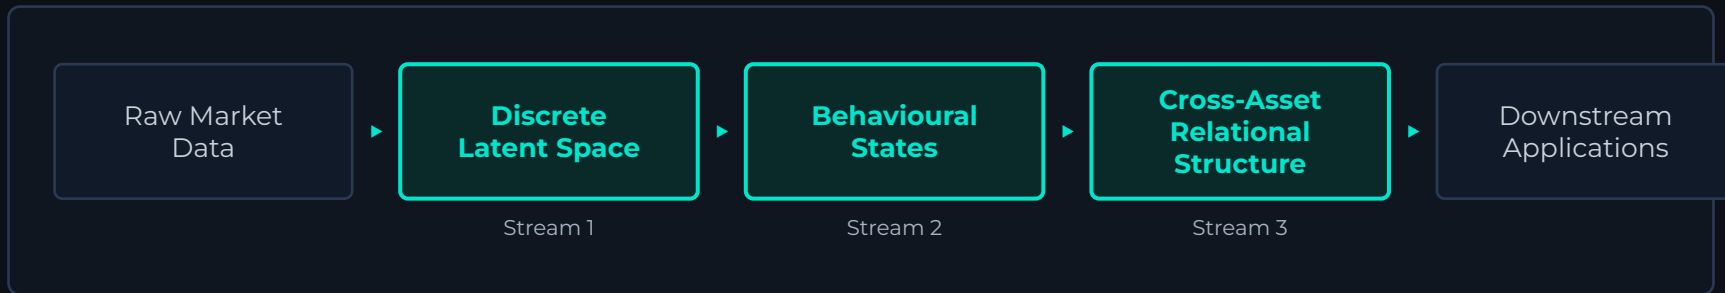
3 Relational Graph Modelling

Geometric Deep Learning

Controlling how the model organises relationships between assets. Cross-asset dependencies and how they rewire as conditions change are learned directly from data. This layer provides the foundation for forward-looking inference.

These are not independent research projects. They form a single integrated system where each layer builds on the one beneath it, producing representations that grow deeper and more general as the architecture matures.

The Foundation Model Principle



The objective is to learn representations of market behaviour at a depth and generality that task-specific models cannot achieve.

A model trained across assets, timescales, and market conditions can learn behavioural transitions and cross-asset propagation patterns that narrow models struggle to infer from a single instrument or a single strategy view. [08]

This is the foundation model principle applied to financial markets: a shared behavioural representation layer that gives downstream models market-wide context. Strategy, execution, and risk modules can specialise on top of the same representations, seeing each instrument in the context of the wider market without rebuilding that context from raw data each time.

Why the Upside Is Asymmetric

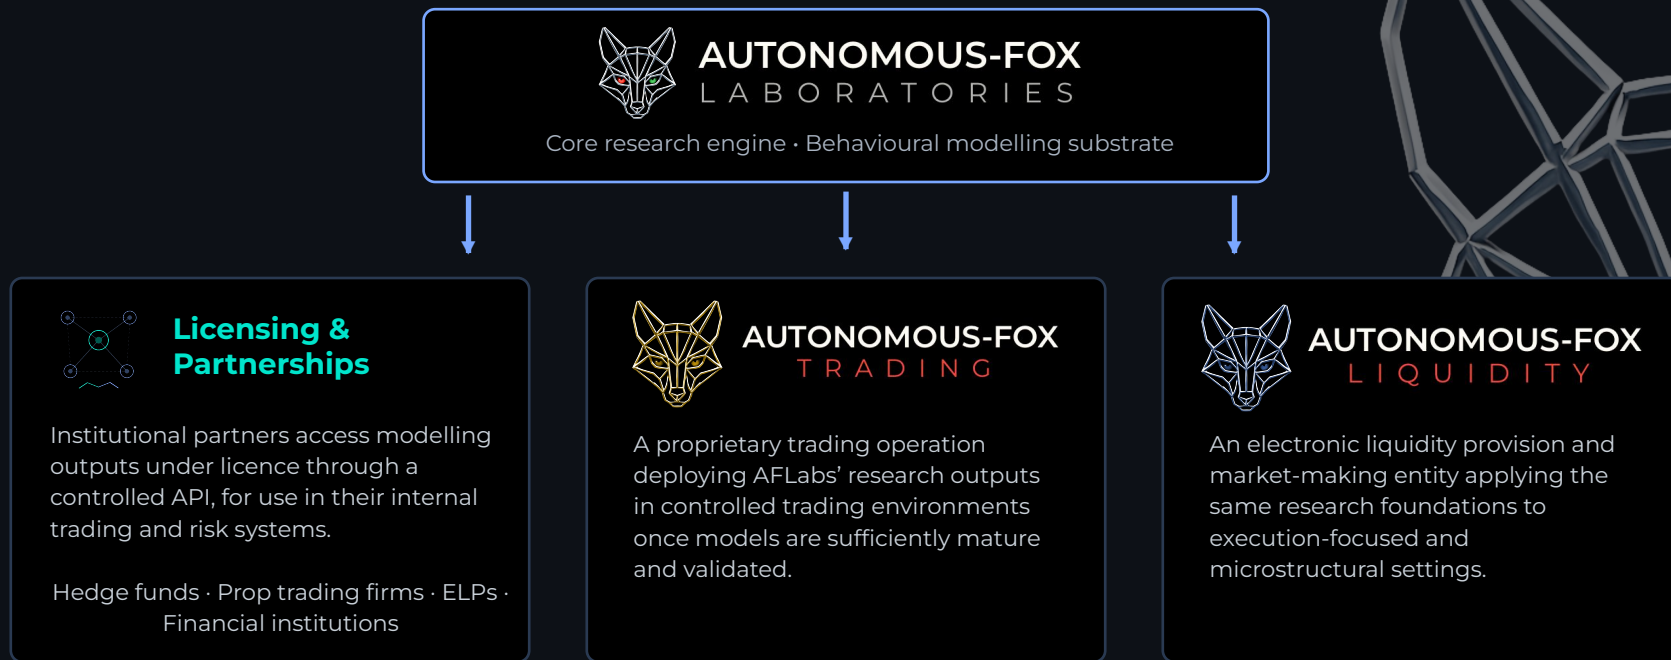
The core investment compounds: improve the substrate once, and every system built on top improves.

Proprietary trading, market making, execution optimisation, risk management, and institutional licensing are all viable downstream pathways from the same research investment.

New applications reuse the same substrate rather than starting from raw data each time.

This kind of leverage is structurally difficult to pursue. It requires a research-first structure with a long time horizon, which is why it has not emerged from within existing trading operations. AFLabs is built from the ground up to pursue it as a primary objective.

Commercial Pathways



At this stage, commercialisation is intentionally secondary to research. The near-term objective is to develop and validate foundational models of market behaviour with clear applicability across quantitative finance.

Risk Profile

The Core Risk

This is a research bet. It is not known whether the information preserved in the learned representations is rich enough to inform downstream trading methods, or whether the learned relationships across assets are sufficient to support forward-looking inference.

The Return Profile

This is not a company that ships a product or doesn't. It is a research programme where the question is how much of the science works, and how far it can be taken.

Improved representations, better inference methods, stronger modelling components — each of those could have standalone commercial value, regardless of whether the full vision is realised.

Risk lies in research difficulty, not execution against a predefined product plan.

Risk in Context

- This is not a greenfield bet. Early work on the discrete tokenisation methodology shows promising results
- The representations suggest economically meaningful structure in live market data
- The remaining challenge is whether preserved information is rich enough, and whether learned relationships are sufficient, to support downstream methods and forward-looking inference

Use of Funds

~80%

of capital directed to R&D

Primary Expenditure

- Senior research scientists and engineers
- Doctoral researchers and research internships
- High-performance computing and cloud infrastructure
- Financial datasets, research premises, IP protection

Funds will not be used for dividends, capital preservation, or low-risk activities. All expenditure is directed toward activities that increase technical capabilities, research velocity, and long-term intellectual property.

Milestones

6 MONTHS

Founding team of five in place across all three research streams

Doctoral studentships secured

12 MONTHS

Core research deployed via spin-out trading entity

Downstream applications tested against live market data

24 MONTHS

Licensing conversations with institutional partners

Headcount and progress review

Team and Stage

Founder

Dr Dimitri Malandreniotis

Principal Research Scientist

Doctorate in machine learning and computational statistics. Twelve years of combined academic and applied experience in machine learning for financial markets, including systematic trading, portfolio construction, statistical arbitrage, and market making.

Growth Trajectory

Scaling up to 12 research scientists within 24 months. Team size grows with funding. Objective: build one of the strongest specialist research teams in behavioural modelling for financial markets.

Hiring the Founding Team

- 2 co-founder level senior research scientists from academic AI groups
- 2 doctoral-level researchers within first 6 months
- **Founding team of 5 across all three research streams**

Academic Collaboration

- UCL Industry Exchange Network (IXN): ~4 master's research projects per year on aligned topics
- Autonomous-Fox Studentships: funded doctoral training at UCL for exceptional candidates
- 2 initial doctoral studentships in discussion with UCL

Tax-Advantaged Investment



Enterprise Investment Scheme (EIS) · Knowledge Intensive Company (KIC)

30%

Income Tax Relief

Investors can claim 30% of the investment against income tax in the year of investment

0%

Capital Gains Tax

Gains on EIS shares are completely free of capital gains tax if held for three years

~61.5%

Effective Loss Relief

Up to 61.5% of the investment recovered through combined income tax and loss relief for additional-rate taxpayers

Knowledge Intensive Company

We expect to qualify for enhanced EIS limits: investors can invest up to £2M per tax year (vs. £1M standard), and the company can raise up to £20M in total EIS funding. KIC status reflects the company's research intensity and deep-technology focus.

Status

HMRC advance assurance application in progress. Expected before end of tax year 2025/26. EIS eligibility is subject to HMRC approval.



AUTONOMOUS-FOX

LABORATORIES

Foundational AI Research for Financial Markets

dimitri@autonomous-fox.ai

+44 7768 596 022

8 Devonshire Square, London, EC2M 4YJ

www.autonomous-fox.ai