AUTONOMOUS-FOX
LABORATORIES

Investment Memorandum

March 2026

1. Who Are We?

Autonomous-Fox Laboratories ("AFLabs") is a UK-based deep-technology research lab conducting foundational artificial intelligence research for financial markets. The company is developing a general-purpose modelling layer for market behaviour: a system designed to learn, update, and maintain coherent asset-specific representations as conditions evolve.

This modelling layer sits beneath application-specific strategies and execution systems, providing shared, reusable representations across alpha generation, portfolio construction, execution, inventory management, and risk control. Rather than constructing isolated models optimised for individual tasks, AFLabs' objective is to build a unified market model whose representations can be specialised across the trading stack without architectural redesign.

The programme builds on twelve years of academic research and applied machine learning development in financial markets, including systematic trading, statistical arbitrage, portfolio construction, and market making. AFLabs extends earlier work on behavioural modelling of market activity using methodological advances that now make deeper and more granular representation learning feasible.

This approach mirrors the role played by foundation models in language and vision, where large-scale representation learning produced reusable modelling infrastructure rather than task-specific tools. No equivalent representation layer currently exists for financial markets. AFLabs is conducting the research required to build one.

The company is research-led, pre-revenue, and IP-focused. Its objective is to create modelling infrastructure whose value compounds as additional strategies, systems, and applications draw on a shared and continuously updating understanding of market behaviour.

2. Why Financial Trading Systems Fail to Scale

Financial markets are among the most challenging real-world systems encountered in applied machine learning. They are non-stationary, with statistical structure that evolves continuously. They are adversarial, adapting to participant behaviour and modelling approaches. They exhibit extremely low signal-to-noise ratios. And they involve complex temporal and cross-asset dependencies that are themselves time-varying and shaped by feedback effects.

Because of this complexity, modern quantitative strategies, proprietary trading, and electronic market-making systems are typically constructed as collections of isolated, task-specific models. Each component learns its own narrow view of the market from raw or lightly processed data, optimised for a single objective and time horizon, with limited awareness of broader market context. These components do not share a common understanding of market behaviour, do not transfer insights between tasks or assets, and must be individually retrained or rebuilt as conditions change.

The result is a modelling architecture that does not scale. As the number of strategies and systems grows, complexity and maintenance costs increase faster than performance. Adaptation occurs through repeated retraining and manual intervention rather than through a coherent, evolving representation of market behaviour: a modelling layer that downstream models can condition on to shape and update their understanding of assets autonomously as market conditions and context change. Different parts of the trading stack therefore operate on incompatible views of the same market, complicating execution, risk management, and system-level coordination.

AFLabs' research is motivated by the view that these limitations are architectural rather than tactical. The core problem is not that existing models are poorly built for their individual tasks, but that the prevailing approach to market modelling, built on narrow, isolated, task-specific components, is fundamentally incapable of producing the kind of deep, transferable understanding of market behaviour that would allow trading systems to scale.

3. Why This Hasn't Been Done Before

Despite the obvious value of a shared, adaptive market model, no major quantitative firm has publicly claimed to have built one. This includes firms with decades of experience, thousands of researchers, and effectively unlimited compute budgets. It is not for lack of ambition or resource. Deep learning, the specific paradigm that enabled general-purpose representation layers in language and vision, has not yet penetrated the core of financial market modelling. Production trading systems at even the most sophisticated firms still overwhelmingly rely on simpler methods: linear models, gradient-boosted trees, classical factor frameworks, and hand-engineered features.

The reasons for this are well understood within the field. All statistical methods struggle with non-stationary data, but simpler approaches survive in practice precisely because their limited capacity constrains how badly they can overfit. Deep networks have the opposite problem: their expressive power, which is the source of their success in other domains, makes them acutely vulnerable in low-signal environments. They memorise noise, learn spurious structure, and collapse when the statistical properties of the data shift. Non-stationarity and low signal content are not incidental features of financial data. They are fundamental properties of the domain, and they explain why deep learning's transformative impact in other fields has not yet been replicated in finance.

What has changed is not the domain, which remains as difficult as it has always been, but the maturation of specific methodological families within machine learning that are better suited to these difficulties. Structured representation learning offers principled approaches to extracting stable abstractions from noisy, non-stationary data. Variational inference methods provide formal frameworks for shaping the structure of learned representations, controlling what the model encodes and how it organises latent space. Graph-based relational modelling allows the system to learn the structure of relationships between assets, how they propagate information through one another, and how that structure rewires itself as market conditions change.

These methodological advances create an opening. But the existence of better-suited tools does not, by itself, solve the problem. Each of these method families was developed for other domains and other purposes. Adapting them to operate reliably on financial data, where distributions shift continuously, signals are orders of magnitude weaker than in language or vision, and the system actively resists being modelled, requires substantial new research. The methods must be fundamentally reworked, not merely applied. Novel architectures must be designed that combine these approaches in ways that have not been attempted. And the resulting systems must be validated under conditions of non-stationarity and adversarial feedback that no existing benchmark captures. This is the work that AFLabs exists to do.

4. What AFLabs Is Building and Why It Is Different

AFLabs is developing a model of financial markets that learns and maintains asset-specific representations of behaviour, each formed in the context of other assets and the wider market environment. These representations evolve as conditions change and are designed to capture structure in market behaviour beyond the reach of task-specific models.

The three research streams described below emerge from twelve years of applied work in financial market modelling. Since 2016, the founder has been developing models that condition on inferred behavioural states of market activity, an approach that showed clear value in live trading environments. However, the classical inference methods available could resolve only a small number of coarse states, imposing hard limits on the depth and granularity of the representations that could be learned. The maturation of the specific methodological families described in Section 3, particularly in structured representation learning and variational inference, is what now makes it feasible to pursue this objective at the scale and resolution the programme targets.

Rather than designing models around predefined regimes, factors, or handcrafted signals, AFLabs' research investigates how market structure can be inferred directly from data in a way that remains coherent as behaviour shifts. The research programme operates on multi-scale, mid-frequency market data spanning timescales from seconds to minutes across multiple asset classes. The programme is organised around three interconnected technical streams, each addressing a distinct open problem in the application of machine learning to financial market systems.

Discrete Tokenisation and Latent Representation Learning

Raw financial time-series are noisy, non-stationary, and lack the intermediate representational structure that has driven progress in language and vision modelling. AFLabs is investigating whether learned vector representations can transform continuous market behaviour into a discrete latent space, where structurally different market behaviours are mapped to distinct, separable representations and noise is suppressed. This is not merely compression. The objective is a representational layer that resolves a far richer and more granular set of behavioural states than existing approaches can distinguish, creating the foundation on which the rest of the modelling architecture is built. Preliminary experimental work supports the feasibility of this direction.

Fine-Grained State Inference

Building on the discrete latent space produced by the tokenisation layer, AFLabs is developing methods for inferring fine-grained state sequences from learned representations. These are not predefined regime categories but behavioural states discovered through learning, at a resolution and granularity that conventional approaches applied to raw market data cannot achieve. The temporal structure of these state sequences is itself learned, capturing how behavioural patterns unfold, and persist over time.

Relational Structure Modelling with Contextual Conditioning

Assets influence one another through economic linkages, capital flows, shared exposures, and contagion dynamics. These relationships shift as market conditions change. AFLabs is investigating graph neural network architectures that learn dynamic relational structure across assets directly from data, rather than relying on pre-specified relationships. The graph model adapts its behaviour based on the latent states inferred by the layer beneath it, so that both the structure of cross-asset relationships and the dynamics operating on that structure respond to the prevailing market environment.

Integration

These three streams form an intentional architectural sequence. The tokenisation layer transforms raw data into a discrete latent space. The state inference layer discovers fine-grained behavioural states within that space. The graph layer models cross-asset dynamics conditioned on those states. Together, they are intended to produce a general-purpose model of market behaviour that can support a wide range of downstream applications.

The objective is to learn representations of market behaviour at a level of depth and generality that no individual task-specific model can construct on its own. A model trained across assets, timescales, and market conditions can capture cross-asset dynamics, regime transitions, and propagation patterns that remain invisible to models designed around a single strategy or asset. This is the foundation model principle: one model whose representations serve many downstream applications, and whose value increases with every advance in the core.

5. Why the Upside Is Asymmetric

Most quantitative trading operations monetise at the strategy layer. In practice, this means building and maintaining collections of models tied to specific strategies, execution tasks, or market settings. These systems therefore tend to scale linearly: each new application requires additional modelling effort, infrastructure, and operational complexity.

A foundational market model would scale differently. A modelling layer that maintains asset-specific, context-aware representations could be reused across multiple strategies, asset classes, and trading functions simultaneously. Additional applications could be built without recreating the underlying modelling effort, allowing value to compound as more systems draw on the same market understanding.

This creates an asymmetric upside. A foundational market model has the potential to support a widening range of downstream applications over time, without requiring the underlying research effort to be rebuilt for each new use case. As additional strategies and systems draw on the same representations, the value of the research base compounds rather than remaining tied to any single application.

This form of leverage is structurally difficult to achieve, which makes it unlikely to be pursued as a deliberate research objective within existing trading firms, even when those firms are large and well capitalised. The constraint is usually not resources, but incentives. Quantitative organisations are typically built to improve strategy performance and deploy models into production, not to invest in long-horizon foundational research whose value may only emerge across many downstream systems over time. Large firms can often afford to continue scaling linearly by adding more strategies, models, and infrastructure, even if that increases fragmentation.

A coherent market model is a cross-cutting research objective that does not fit naturally inside that structure. AFLabs is structured specifically to pursue this opportunity.

6. The Behavioural Modelling Layer

The preceding section described the asymmetric economics of a foundational market model. This section describes what that model concretely delivers to any system built on top of it.

For each asset, the system maintains a learned behavioural state: a representation of how that asset is currently behaving, inferred directly from data. These are not price levels, returns, or conventional indicators. They are learned representations that capture the behavioural mode the asset is in, how long that mode has persisted, and whether it is stable or in transition. The objective is to resolve behavioural structure at a depth and granularity that conventional methods applied to raw market data cannot achieve.

Each asset's behavioural state is formed in the context of every other asset in the system. The relational graph layer described in Section 4 is designed to learn which assets are influencing one another at any given time, how those relationships are structured, and how they are changing. The representation of any single asset therefore reflects not only its own behaviour but the wider environment around it: which relationships are currently active, where stress or momentum is propagating, and how the behaviour of related assets is shaping what that asset is doing.

The system also produces forward-looking trajectory inference: how the behavioural state of each asset is expected to evolve. This is not a forecast conditioned on the asset's own history alone. It is conditioned on the current behavioural modes of related assets, the paths by which those assets arrived at their current states, and the relational structure connecting them at that moment. The expected trajectory for any single asset is therefore shaped by the wider observable behavioural landscape across the system.

That is what downstream models receive. Instead of each model building its own narrowly scoped view of the market from raw data, it conditions on a shared layer that provides awareness of what behavioural mode each asset is in, how those modes are evolving across related assets, what the current relational structure looks like, and how the system expects behaviour to change next.

The consequence is that each downstream model inherits a deep, continuously updated understanding of market behaviour that it could not construct from raw data alone. This allows the trading stack as a whole to operate with a level of market awareness that isolated, task-specific models struggle to achieve in a coherent way, with the further practical consequence that downstream models require significantly less retraining to stay aligned with changing market context.

Early results provide confidence to extend the research programme to the full architecture described in Section 4. The core tokenisation methodology transforms continuous market behaviour into discrete representations that recover semantic structure in market activity and directly address the noise problem that has historically constrained the application of deep learning to financial data. Modern sequence models applied to those representations have proven effective at modelling and inferring behavioural state trajectories. These results establish a credible foundation for the first two research streams described in Section 4. The relational graph modelling stream, which extends these representations across assets, remains the more prospective element of the programme.

7. Commercial Pathways

AFLabs is a research laboratory. It does not trade, does not build trading strategies, and does not operate in financial markets. Its commercial model is the creation of foundational intellectual property and the operation of the core modelling infrastructure that the IP produces. AFLabs holds the IP, operates the models, and licenses access to the live inference stream they generate.

What clients receive is not static technology or a frozen model. They receive a continuously updating stream of behavioural representations, state trajectories, and relational structure: the outputs described in Section 6, delivered as a live service. Clients use this stream to condition their own downstream models and trading systems. Because the inference stream depends on the models continuing to run and update, the commercial relationship is inherently recurring rather than transactional. Each client depends on sustained access to a system that only AFLabs operates.

The licensing model applies uniformly across all clients. Whether the licensee is a purpose-built downstream entity or an institutional partner, the commercial relationship is the same: AFLabs provides access to the inference stream on contracted commercial terms, the client deploys it within their own operational context, and AFLabs retains full ownership of the underlying intellectual property and infrastructure.

The first two licensees will be purpose-built entities established specifically to deploy AFLabs' research in operational settings:

- Autonomous-Fox Trading: a separately incorporated proprietary trading operation that will license access to AFLabs' inference stream and deploy it within controlled trading environments.
- Autonomous-Fox Liquidity: a separately incorporated electronic liquidity provision and market-making entity that will license the same stream for execution-focused and microstructural settings.

These entities are not subsidiaries, affiliates, or divisions of AFLabs. They are independent companies whose sole relationship to the lab is as contracted licensees. AFLabs will have no operational involvement in their trading activity, strategy development, or market participation. They bear their own operational, regulatory, and market risk entirely independently of the lab.

These first licensees serve a dual purpose beyond revenue. They provide a live deployment environment in which AFLabs' models are tested against real market conditions and real trading applications. This feedback loop is the most cost-effective mechanism available for validating and hardening the intellectual property: it allows the lab to observe how its research performs under genuine operational stress without the lab itself needing to build or manage trading systems. Early deployment through dedicated licensees therefore accelerates the research while simultaneously establishing the commercial model.

As the technology matures, the same licensing model extends to institutional clients: quantitative hedge funds, proprietary trading firms, electronic market makers, and financial technology partners. Because the modelling infrastructure is general-purpose, it can support a range of client-specific applications without requiring AFLabs to operate in those domains directly. Research and deployment partnerships represent a related pathway, allowing AFLabs to collaborate on applied problems while retaining ownership of the underlying intellectual property.

This structure preserves both focus and optionality. The downstream entities are the first clients. Institutional partners and licensees follow as the research matures. The value of what investors hold — the intellectual property and modelling infrastructure inside AFLabs — compounds with every client that draws on it.

8. Research Risk and Return Profile

AFLabs is pursuing a research problem that is technically difficult because it cannot be solved through incremental refinement of existing trading models. It is not known whether stable, reusable representations of market behaviour can be learned from non-stationary, low-signal data in a manner that generalises across assets, time periods, and evolving market conditions. The discrete latent representations at the core of the programme may fail to preserve sufficient structural information, the inferred behavioural states may not correspond to economically meaningful distinctions, or the relational graph models may not learn useful cross-asset structure from unlabelled data.

This uncertainty defines the risk profile of the company. Progress may be uneven, and some research directions may fail to generalise or scale. However, the programme is deliberately structured so that partial technical advances, such as improved representations, inference methods, or modelling components, can still have standalone value and inform multiple downstream applications. Importantly, early experimental results have already demonstrated that the core tokenisation methodology can recover meaningful behavioural structure from raw market data, providing direct evidence that the foundational layer of the programme is technically viable. This de-risks the first research stream and provides a validated base on which the more prospective elements of the architecture are built.

For investors, the risk lies in the difficulty of the research itself rather than in execution against a predefined product plan. The potential reward arises from the leverage inherent in reusable market representations, where even limited success can support multiple strategies, systems, and commercial pathways.

9. Use of Funds

Capital raised under EIS will be used exclusively to support AFLabs' research programme and the development of its foundational modelling infrastructure. The company expects approximately 80% of total expenditure to be directed to research and development activity.

The largest category of expenditure is employment of research scientists and engineers, including senior research appointments, doctoral researchers, and research internship programmes conducted in collaboration with academic institutions. Additional uses include high-performance computing and cloud infrastructure, acquisition and processing of large-scale financial datasets, research premises, and the creation and protection of intellectual property.

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

The purpose of this funding is to establish the company's core research capability and foundational modelling infrastructure in the near term. An initial downstream deployment pathway is expected to become active at around 12 months, when a separately incorporated quantitative trading entity, operating as a contracted licensee of AFLabs' intellectual property, will begin deploying downstream models built on top of the core research. This deployment will provide an essential source of real-world validation and feedback, allowing the lab to observe how its models perform under live market conditions and to refine the research accordingly. Broader commercialisation through institutional licensing and partnerships is expected to begin from around 24 months.

10. Team

AFLabs is founded and led by Dr Dimitri Malandreniotis, Principal Research Scientist. Dr Malandreniotis brings over twelve years of experience at the intersection of machine learning research and financial market application, spanning systematic trading, portfolio construction, statistical arbitrage, and market making, alongside doctoral-level research in machine learning and computational statistics. This combination of deep theoretical training and extensive applied experience in the specific domain AFLabs operates in is central to the company's research programme: the work requires not only methodological capability but sustained, first-hand understanding of the environment in which these methods must operate.

The company is building a research-intensive team designed to grow in step with available funding and research progress. Immediate hiring priorities include two senior research scientists at co-founder level, who will share responsibility for the direction and execution of the core research programme. The company is in active discussions with multiple candidates from leading academic AI research groups, with conversations progressing at different stages. Within the first six months, the company intends to recruit two doctoral-level researchers to work directly on the research streams described in Section 4. This initial team of five establishes the core research capability required to advance all three streams in parallel.

Beyond the founding team, AFLabs has established a structured pipeline for identifying and developing research talent through its collaboration with University College London. The company participates annually in the UCL Industry Exchange Network (IXN) programme, supervising up to four master's-level research projects per year on topics aligned with the company's research objectives. High-performing students may be offered positions as entry-level researchers within AFLabs or, where appropriate, considered for doctoral training through Autonomous-Fox funded Studentships. The company is in early discussion with UCL regarding the establishment of this studentship pathway and hopes over time to fund up to two doctoral candidates per year.

This talent model is designed to create compounding research capacity over time. The founding team provides immediate depth and technical leadership. The UCL collaboration generates a recurring intake of researchers tested against real problems before any hiring commitment is made. The funded studentships create a longer-term pipeline of scientists trained specifically on the problems that define AFLabs' programme. The result is a team-building approach that scales with the research, rather than ahead of it.

Summary for Investors

AFLabs is a deep-technology research company addressing a foundational gap in how financial markets are modelled. While modern trading and market-making systems rely on increasingly complex collections of task-specific models, there is no general-purpose model of market behaviour that adapts as conditions change and can be specialised across strategies, execution systems, and risk frameworks.

Deep learning, the paradigm that enabled general-purpose representation layers in language and vision, has not yet had an equivalent impact in financial markets, despite the resources available to the largest quantitative firms. Recent advances in self-supervised representation learning, sequence modelling, and geometric deep learning have created an opening, but adapting these methods to operate reliably on financial market data requires fundamental new research.

AFLabs is conducting that research. The work is organised around three interconnected research streams: discrete tokenisation of market behaviour, fine-grained state inference over learned representations, and relational graph modelling with contextual conditioning. Together, these are intended to produce a continuously updating modelling layer whose live inference stream — a real-time feed of behavioural states, trajectory forecasts, and relational structure — downstream trading, execution, and risk systems can condition on, rather than forcing each model to build and rebuild its own isolated view of the market.

The intended outcome is a fundamentally different operating model for trading systems: one in which new strategies and applications can be developed more rapidly, adaptation occurs at the level of the underlying model rather than through repeated rebuilds, and complexity grows more slowly than capability. This shifts value creation from incremental performance gains to infrastructure-level leverage.

The company is research-led, pre-revenue, and IP-focused. Its commercial model centres on the creation of foundational intellectual property and the licensing of the live inference stream that the modelling infrastructure produces, with revenue generated through contracted licensing to downstream entities and, over time, to institutional clients. AFLabs expects to qualify for EIS and Knowledge-Intensive Company relief, providing additional downside protection for qualifying investors.

