Adaptive Market Intelligence

A Probabilistic Multi-Agent Architecture for Real-Time Market-State Inference, Strategy Selection and Risk-Governed Financial Decision-Making

PhD Research Proposal Outline

Candidate: David Maguire

Date: 4 August 2026

Updated working document for supervisor discussion and proposal development

CENTRAL PROPOSITION

Financial-market systems should first infer what type of market is currently being observed, quantify uncertainty and identify suitable strategy families. They should not move directly from raw data to directional trading decisions.

Proposal at a Glance

Field	Updated proposal
Working title	Adaptive Market Intelligence: A Probabilistic Multi-Agent Architecture for Real-Time Market-State Inference, Strategy Selection and Risk-Governed Financial Decision-Making
Academic alternative	Probabilistic Multi-Agent Market-State Inference and Adaptive Strategy Selection under Non-Stationarity and Uncertainty
Candidate	David Maguire
Primary domain	Quantitative finance, machine learning, financial econometrics and agentic decision systems
Primary contribution	A hierarchical architecture that converts heterogeneous real-time data into calibrated specialist evidence, competing market-state hypotheses, consensus estimates and regime-conditioned strategy suitability
Initial empirical scope	NASDAQ-100 ecosystem using NQ futures, QQQ, volatility, rates, breadth, liquidity, cross-asset and macro-event information
Core methods	Regime detection, change-point detection, probabilistic calibration, ensemble learning, financial NLP, constrained agent coordination, strategy selection and formal risk control
Commercial pathway	Institutional market-intelligence and decision-support platform, with later extensions into allocation, risk overlays and execution support

Executive Summary

This research proposes a real-time Adaptive Market Intelligence architecture designed to infer the current financial-market state before selecting a strategy or action. The central problem is not framed as direct price prediction. It is framed as point-in-time inference under non-stationarity: what kind of market is currently being observed, how certain is that assessment, which evidence supports or contradicts it, and which strategy families are appropriate under those conditions?

The architecture would ingest synchronised market, microstructure, volatility, cross-asset, macroeconomic and textual data. Specialist quantitative models would transform those inputs into bounded and calibrated evidence concerning trend, volatility, liquidity, structural change, macro drivers and cross-asset confirmation. A multi-agent reasoning layer would then compare competing market-state hypotheses rather than simply average incompatible model scores.

The resulting market-state representation would be multidimensional and probabilistic. It could identify, for example, a macro-driven risk-off transition with expanding volatility, weakened but orderly liquidity and elevated uncertainty. A separate strategy council would assess the suitability of trend, mean-reversion, long-volatility, statistical-arbitrage, market-neutral and cash-preservation strategies. A deterministic risk agent would retain final authority to scale, reject or restrict proposed actions.

The academic contribution would be the design and rigorous evaluation of a calibrated, hierarchical and auditable market-state inference system. The commercial contribution would be a blueprint for an institutional market-intelligence platform that provides real-time regime context, transition alerts, strategy suitability, risk recommendations, evidence traceability and decision history.

What Changed from the Previous Outline

Previous emphasis	Updated emphasis
One principal Market Regime Agent	Multiple specialist evidence models covering trend, volatility, liquidity, macro narrative, cross-asset confirmation and structural change
Combined state passed directly to an RL Trading Agent	Explicit hypothesis-generation and consensus layer before any strategy decision
Primary downstream output was exposure or trade action	Primary output is probabilistic market-state intelligence and strategy suitability
RL central from an early stage	RL is optional and introduced only after regime inference and deterministic allocation baselines are validated
Agents loosely described as functional modules	Clear distinction between quantitative models, specialist agents, consensus mechanism, strategy agents and deterministic risk control
Commercial value implied	Commercial product, user outputs, proprietary data assets and staged path to market made explicit

1. Introduction and Background

Financial markets are dynamic systems in which relationships between returns, volatility, liquidity, order flow, interest rates, economic information and investor behaviour vary through time. A forecasting rule that is effective during a liquid directional trend may become unreliable during an event-driven repricing, volatility shock or liquidity withdrawal. The environment is therefore non-stationary and only partially observable.

Many financial-machine-learning studies nevertheless frame the central problem as predicting the next return or selecting buy, sell and hold actions. This can lead to systems that appear successful in historical samples but fail when the underlying market process changes. The proposed research takes a different position: before a model selects a strategy, it should estimate the state of the decision environment and the uncertainty attached to that estimate.

The market state is not assumed to be a single label such as bull, bear or sideways. It is represented through simultaneous dimensions, including directionality, volatility, liquidity, macro influence, cross-asset confirmation, structural change and transition risk. This better reflects the fact that a market may be trending while volatility expands, liquidity weakens and macro information dominates price formation.

- Directional condition: upward trend, downward trend, neutral or unstable.
- Volatility condition: contracting, normal, expanding or dislocated.
- Liquidity condition: healthy, weakening, stressed or impaired.
- Dominant driver: technical, macroeconomic, corporate-event, positioning or unknown.
- Transition condition: stable, emerging, changing or structurally uncertain.
- Cross-asset confirmation: confirming, mixed or contradicting the local market signal.

2. Research Problem

1. Static models continue to apply historical relationships after the environment has changed.
2. Single-label regime classifications oversimplify overlapping market conditions.
3. Heterogeneous numerical and textual information is often processed separately and reconciled informally.
4. Model confidence is frequently uncalibrated, so a stated probability cannot be interpreted as a reliable probability.

5. Ensembles commonly average outputs even when models answer different questions or operate over different horizons.
6. Trading systems move too quickly from observations to actions without explicit market-state hypothesis testing.
7. Reinforcement learning is often tasked with direct trading before the state representation has been shown to be stable or informative.
8. Commercial systems require auditability, latency control, deterministic risk limits and evidence traceability, not only backtest returns.

CENTRAL RESEARCH PROBLEM

Can a hierarchical system of specialist quantitative models and bounded reasoning agents infer multidimensional financial-market states and transitions in real time, with calibrated uncertainty, and can those inferences improve out-of-sample strategy selection and risk management relative to static, regime-unaware and single-model alternatives?

3. Research Aim and Objectives

3.1 Overall Research Aim

To design, develop and evaluate a real-time Adaptive Market Intelligence framework that converts heterogeneous point-in-time financial data into calibrated specialist evidence, reconciles competing market-state hypotheses, estimates multidimensional market conditions and supports regime-conditioned strategy selection under deterministic risk constraints.

3.2 Research Objectives

Objective	Description
O1	Define an economically interpretable multidimensional taxonomy of market states and transitions.
O2	Construct a timestamp-safe data architecture combining market, microstructure, volatility, cross-asset, macroeconomic and textual information.
O3	Develop specialist model families for trend, volatility, liquidity, structural change, macro narrative and cross-asset confirmation.
O4	Calibrate specialist probabilities and estimate reliability conditional on data quality, horizon and previously observed market conditions.
O5	Develop a bounded multi-agent consensus mechanism that compares competing market-state hypotheses and records supporting and contradictory evidence.
O6	Evaluate whether inferred market states improve strategy suitability, exposure management and risk outcomes relative to regime-unaware alternatives.
O7	Test whether adaptive expert weighting provides incremental value over transparent static weighting and Bayesian or ensemble baselines.
O8	Design an auditable decision and review framework capable of tracing each output from source data to models, consensus, strategy recommendation and realised outcome.

Objective	Description
O9	Assess the architecture as the technical and intellectual-property foundation for an institutional market-intelligence platform.

4. Research Questions and Hypotheses

4.1 Principal Research Question

PRINCIPAL RESEARCH QUESTION
Can calibrated specialist models and a bounded multi-agent consensus architecture identify multidimensional market states and regime transitions in real time, and does this information improve out-of-sample strategy selection, risk control and decision robustness?

4.2 Supporting Research Questions

Ref.	Research question
RQ1	Which statistical, econometric and machine-learning methods most reliably identify trend, volatility, liquidity and structural-change states out of sample?
RQ2	Can overlapping market dimensions be estimated more reliably than a single mutually exclusive regime label?
RQ3	How should probabilities from heterogeneous models be calibrated, normalised and compared across different horizons?
RQ4	Can structured information extracted from economic announcements and financial text improve identification of market drivers beyond numerical market data?
RQ5	Does explicit comparison of competing market-state hypotheses outperform simple averaging, stacking or voting ensembles?
RQ6	Does market-state information improve the selection, activation and risk scaling of strategy families after realistic costs?
RQ7	Can expert reliability be estimated conditionally, such that models receive different weights in different market environments?
RQ8	Are adaptive weighting or reinforcement-learning methods superior to transparent deterministic and Bayesian weighting methods?
RQ9	How stable, explainable and auditable are the system outputs across assets, periods, latency assumptions and repeated trials?
RQ10	Which components create commercially useful decision intelligence even if direct trading performance is not consistently superior?

4.3 Proposed Hypotheses

Ref.	Hypothesis
H1	A multidimensional market-state representation will exhibit greater temporal stability and downstream usefulness than a single-label regime classification.
H2	Calibrated ensemble probabilities will provide more reliable out-of-sample state estimates than uncalibrated model confidence scores.
H3	Structured macro-event and narrative features will add incremental value primarily during event-driven and transition periods.
H4	Explicit hypothesis reconciliation will improve state classification and uncertainty estimation compared with simple majority voting or averaging.
H5	Regime-conditioned strategy selection will improve risk-adjusted performance, drawdown control or tail-risk behaviour relative to static strategy allocation.
H6	The greatest incremental value will arise during regime transitions and periods of disagreement between specialist models.
H7	Condition-specific expert weighting will outperform equal weighting, but complex adaptive methods will not necessarily outperform well-designed transparent baselines.
H8	Deterministic risk controls will improve robustness and practical viability even where they reduce gross return.

5. Proposed Conceptual Architecture

5.1 Hierarchical Decision Process

Layer	Function	Illustrative output
Layer 1	Data acquisition	Real-time and historical feeds with source, event and available-at timestamps.
Layer 2	Point-in-time data platform	Raw, cleaned and research-ready layers; synchronisation, quality flags and feature lineage.
Layer 3	Specialist quantitative models	Trend, volatility, liquidity, structural change, macro narrative and cross-asset evidence.
Layer 4	Evidence standardisation	Calibrated probability, reliability, horizon, freshness, model version, supporting evidence and anomaly warnings.
Layer 5	Market-state hypothesis engine	Construction and comparison of competing explanations of the current market.
Layer 6	Consensus and uncertainty	Posterior market-state estimate, disagreement, transition probability and confidence.

Layer	Function	Illustrative output
Layer 7	Strategy council	Suitability assessment for trend, mean reversion, volatility, statistical arbitrage, market neutral and cash preservation.
Layer 8	Risk governance	Deterministic exposure, leverage, liquidity, drawdown, concentration and execution constraints.
Layer 9	Outcome review and learning	State accuracy, strategy suitability, model contribution, overrides and realised outcomes.

5.2 Distinction Between Models and Agents

The proposal distinguishes a quantitative model from an agent. A model estimates a defined statistical quantity. An agent owns a bounded task, may consult one or more models, applies validation rules and communicates through a fixed schema. This distinction is necessary to avoid using agent terminology as a substitute for methodological precision.

Component	Function	Example
Quantitative model	Estimates a statistical state or quantity	EGARCH estimates conditional volatility; HMM estimates latent state probabilities.
Specialist agent	Validates and summarises related model evidence	Volatility Specialist compares EGARCH, realised-volatility and anomaly-model outputs.
Consensus agent	Compares competing market-state hypotheses	Distinguishes an ordinary technical pullback from a macro-driven risk-off transition.
Strategy agent	Assesses suitability, not raw direction alone	Mean-Reversion Agent recommends inactivity during event-driven volatility expansion.
Risk agent	Applies non-negotiable constraints	Reduces exposure when spreads widen or drawdown limits are approached.

5.3 Specialist Evidence Domains

Specialist	Primary data	Candidate methods	Structured output
Trend and persistence	Returns, breadth, VWAP, momentum, cross-sectional participation	HMM, Kalman filter, state-space model, gradient boosting	Direction, persistence, transition probability and model agreement
Volatility	Realised and implied volatility, ranges, term structure, volatility-of-volatility	GARCH family, HAR-RV, boosting, sequence models	Expansion probability, expected horizon, tail-risk condition
Liquidity and microstructure	Spreads, depth, order imbalance, trade intensity, price impact	Anomaly detection, supervised stress model, order-book features	Liquidity condition, execution risk, stress probability
Structural change	Distribution shifts, correlation changes, feature drift	Bayesian online change-point detection, CUSUM, drift detection	Break probability, detection delay, affected variables

Specialist	Primary data	Candidate methods	Structured output
Macro and narrative	Announcements, central-bank communication, financial news	Schema-constrained LLM, event extraction, surprise measurement	Event class, direction, novelty, duration, affected assets and evidence
Cross-asset confirmation	Rates, FX, volatility, sectors, credit, commodities	Dynamic correlation, factor models, graphical methods	Confirming and contradicting markets; risk-on or risk-off evidence

6. Data Architecture and Real-Time Example

6.1 Initial Empirical Market

The initial prototype would focus on the NASDAQ-100 ecosystem. This is sufficiently rich to test macro sensitivity, technology-sector concentration, volatility, rates and cross-asset interaction while remaining narrow enough for a three-year doctoral programme.

- NQ futures and QQQ ETF prices, trades and quotes.
- NASDAQ-100 constituent returns and market breadth.
- VIX or VXN and selected options-derived volatility measures.
- Two-year and ten-year US Treasury yields or liquid proxies.
- US dollar and selected cross-asset risk indicators.
- Scheduled US macroeconomic releases and central-bank events.
- Timestamped financial-news or announcement data where licensing permits.

6.2 Data-Layer Design

Layer	Purpose	Contents
Bronze	Preserve source evidence	Unmodified vendor messages, source timestamp, receive timestamp, identifiers and data-quality flags.
Silver	Create a clean point-in-time record	Deduplicated, normalised and time-aligned observations with corporate-action and session handling.
Gold	Support reproducible modelling	Feature vectors, labels, model outputs, calibration data, state estimates and decision records.

6.3 Example One-Minute Decision Record

Feature	Illustrative value
Decision timestamp	14:31:00 UTC
NQ one-minute return	-0.18%
NQ fifteen-minute return	-0.62%
Thirty-minute realised volatility	24.7% annualised
Bid-ask spread z-score	2.1
Order imbalance	-0.43
NASDAQ breadth	31% advancing
Volatility-index change	+4.2%

Feature	Illustrative value
Two-year yield change	+6.3 basis points
Macro-event classification	Hawkish inflation surprise
Minutes since release	1

6.4 Example Specialist Evidence

Specialist	State	Probability	Important evidence
Trend	Downward trend	0.78	Negative returns, weak breadth and price below VWAP
Volatility	Volatility expansion	0.91	Realised-volatility jump and volatility-index confirmation
Liquidity	Deteriorating	0.74	Wider spreads, lower depth and negative imbalance
Macro narrative	Hawkish rate repricing	0.84	Inflation surprise and rapid yield adjustment
Cross-asset	Risk-off confirmation	0.69	Yields, dollar and volatility broadly confirm

6.5 Competing Market-State Hypotheses

Hypothesis	Supporting evidence	Contradictory evidence	Posterior
Ordinary technical pullback	Prior overextension and slowing momentum	Move began at macro release; yields and volatility reacted simultaneously	18%
Macro-driven risk-off transition	Inflation surprise, rate repricing, weak breadth, volatility expansion and liquidity deterioration	Credit stress remains limited	76%
Full liquidity crisis	Wider spreads and reduced depth	Market remains orderly and credit evidence is absent	6%

ILLUSTRATIVE FINAL MARKET STATE

Primary state: macro-driven risk-off transition (76%). Secondary state: volatility expansion (83%). Transition probability: 71%. Liquidity: weakened but orderly. Expected horizon: several hours to several days. Overall confidence: 73%.

7. Strategy Suitability and Risk Governance

7.1 Strategy Council

The strategy layer is intentionally separated from market-state inference. Strategy agents receive the same consensus state and independently assess whether their strategy should be active,

inactive or risk-reduced. This separation makes it possible to test whether market-state intelligence is valuable without claiming that one trading model can forecast every environment.

Strategy family	Illustrative suitability	Illustrative interpretation
Directional trend	82%	Favourable because direction, macro driver and cross-asset evidence are aligned.
Long volatility	71%	Potentially suitable, although entry cost must reflect already elevated implied volatility.
Market neutral	54%	Moderate suitability where idiosyncratic dispersion remains available.
Statistical arbitrage	39%	Reduced confidence because correlations and spreads may be unstable.
Mean reversion	17%	Inactive because event-driven repricing and volatility expansion increase continuation risk.

7.2 Deterministic Risk Authority

- Maximum gross and net exposure.
- Position-size limits conditional on volatility and liquidity.
- Leverage, concentration and correlation limits.
- Maximum drawdown and Expected Shortfall controls.
- Minimum liquidity and maximum spread thresholds.
- Scheduled-event and overnight restrictions.
- Transaction-cost and slippage assumptions.
- Mandatory no-action state when evidence is stale, contradictory or outside the training distribution.

8. Theoretical and Conceptual Foundations

Foundation	Relevance
Adaptive Markets Hypothesis	Supports the proposition that market efficiency and strategy performance vary with competitive and environmental conditions.
Market microstructure	Explains liquidity, order flow, price discovery and execution risk.
Latent-state and regime-switching models	Provide formal methods for inferring unobserved market conditions.
Bayesian inference and hypothesis comparison	Support posterior updating and explicit comparison of competing explanations.
Ensemble learning and expert advice	Provide baselines for model combination and condition-specific weighting.
Probability calibration and uncertainty estimation	Allow confidence scores to be interpreted and tested.
Partially observable decision processes	Represent the fact that the true market state is not directly observed.
Multi-agent systems	Support bounded roles, communication, conflict resolution and attribution.

Foundation	Relevance
Financial NLP and LLMs	Transform unstructured events and narratives into structured, evidence-linked features.
Model risk and AI governance	Provide requirements for traceability, reproducibility, override and human oversight.

9. Proposed Methodology

9.1 Research Design

The project would follow a staged empirical design. Each component would first be evaluated independently. Integration would proceed only where ablation evidence demonstrates incremental value. The research would prioritise point-in-time integrity, transparent baselines and out-of-sample performance over architectural complexity.

9.2 Study Structure

Study	Primary question	Expected output
Study 1: Market-state taxonomy and detection	Can multidimensional states and transitions be identified consistently?	Regime taxonomy, labelled evaluation framework and benchmark comparison.
Study 2: Calibrated specialist evidence	Can heterogeneous model probabilities be calibrated and conditionally weighted?	Evidence schema, reliability model and ensemble baselines.
Study 3: Macro and narrative contribution	Does structured text improve event-driven state inference?	Timestamp-safe event features and incremental-value analysis.
Study 4: Consensus and hypothesis reconciliation	Does explicit hypothesis comparison outperform averaging and voting?	Bounded consensus mechanism and disagreement measures.
Study 5: Strategy suitability and risk	Does state information improve strategy activation and exposure control?	Regime-conditioned strategy framework with realistic costs.
Study 6: Integrated system and governance	Is the full system robust, explainable and commercially useful?	End-to-end prototype, ablation study, audit trail and product blueprint.

9.3 Baselines

- Rule-based volatility, trend and liquidity states.
- Hidden Markov and Markov-switching models.
- K-means and Gaussian mixture models.
- Bayesian online change-point detection.
- Logistic regression, Random Forest and gradient boosting.
- Equal-weight, majority-vote, Bayesian model averaging and stacking ensembles.
- Static strategy allocation and transparent regime-conditioned allocation.
- Single-agent architecture and regime-unaware strategy models.

9.4 Validation Design

- Chronological train, validation and test periods.
- Walk-forward and rolling-window evaluation.
- Purged cross-validation and embargo where labels overlap.
- Available-at timestamps and frozen historical information sets.

- Transaction costs, bid-ask spreads, slippage and delayed execution.
- Regime-specific, transition-specific and event-specific evaluation.
- Repeated random seeds and sensitivity analysis.
- Stress testing across different monetary-policy and volatility environments.
- Out-of-distribution and data-staleness testing.

9.5 Role of Reinforcement Learning

Reinforcement learning would not be assumed to be necessary. It would be tested only after the state representation, specialist evidence and deterministic strategy mappings have been validated. Candidate RL applications include expert weighting, strategy allocation, exposure scaling or conflict resolution. It would be compared directly with transparent rules, contextual bandits, Bayesian updating and supervised meta-models. A finding that simpler methods are superior would remain an academically valid result.

10. Evaluation Framework

Evaluation area	Measures
State identification	Accuracy where labels exist, adjusted Rand index, persistence, stability, detection delay and economic interpretability.
Probabilistic quality	Log loss, Brier score, expected calibration error, reliability diagrams and sharpness.
Transition detection	Lead or delay, false alarms, missed transitions and time-to-detection.
Consensus quality	Hypothesis ranking, disagreement resolution, contradiction handling and confidence degradation.
Strategy suitability	Return, Sharpe, Sortino, Calmar, drawdown, Expected Shortfall, turnover and tail behaviour.
Execution realism	Spread, slippage, market impact, latency and implementation shortfall.
Agent contribution	Ablation, Shapley-style contribution where appropriate, overrides and incremental value.
Robustness	Block bootstrap intervals, multiple-testing controls, Deflated Sharpe Ratio and Probability of Backtest Overfitting.
Governance	Traceability, reproducibility, schema compliance, source attribution and deterministic override frequency.
Commercial usefulness	Alert timeliness, interpretability, user actionability, stability and value without direct execution.

11. Expected Contribution and Significance

11.1 Methodological Contribution

- A hierarchical distinction between data, quantitative models, specialist agents, consensus reasoning, strategy suitability and risk authority.
- A multidimensional probabilistic representation of market state rather than a single mutually exclusive regime label.
- A common evidence schema containing calibrated probability, conditional reliability, horizon, freshness, support, contradiction and distribution-shift information.

- A bounded hypothesis-reconciliation mechanism for combining heterogeneous financial evidence.
- A rigorous comparison of adaptive expert weighting against transparent ensemble baselines.

11.2 Empirical Contribution

- Evidence on which regime dimensions are detectable and useful out of sample.
- Evidence on whether market-state information improves strategy selection after realistic costs.
- Evidence on when textual and macro information adds value beyond market data.
- Evidence on whether complexity is most useful during transitions and specialist disagreement.

11.3 Practical and Governance Contribution

- A real-time decision architecture that can produce useful intelligence without granting an LLM trading authority.
- An auditable path from raw observation to model evidence, consensus state, strategy recommendation and realised outcome.
- A framework for deterministic risk controls, no-action states and human oversight.

12. Commercialisation Blueprint

12.1 Initial Product Proposition

The first commercial product should be a decision-support platform rather than an autonomous hedge fund or signal-selling service. It would provide institutional users with a live interpretation of market conditions, evidence and uncertainty.

Product output	Illustrative user value
Current market state	A concise description of trend, volatility, liquidity, dominant driver and transition risk.
Confidence and disagreement	Shows whether evidence is aligned or contested.
Transition alerts	Identifies when the market may be leaving a previously stable state.
Strategy suitability	Shows which strategy families are historically compatible or incompatible with the inferred state.
Risk recommendation	Suggests exposure reduction, position-size changes or no-action conditions.
Evidence trace	Explains which market, macro and textual observations drove the conclusion.
Historical analogues	Retrieves comparable prior states and subsequent outcomes without claiming deterministic repetition.

12.2 Potential Customers

- Asset managers and hedge funds.
- Bank treasury, trading and risk functions.
- Family offices and proprietary trading firms.
- Execution and electronic-trading desks.
- Wealth-management and investment-platform providers.
- Risk, model-validation and market-surveillance teams.

12.3 Defensible Intellectual Property

- Proprietary point-in-time data engineering and feature lineage.
- Market-state taxonomy and transition definitions.

- Calibration and conditional expert-reliability histories.
- Hypothesis-reconciliation logic and disagreement measures.
- Regime-to-strategy suitability mappings.
- Historical analogue and outcome library.
- Decision, override and attribution audit trail.
- Institution-specific tuning and governance controls.

12.4 Staged Route to Market

Stage	Research or product outcome
Stage 1	Research dashboard using delayed or historical data; validates state inference and explanation.
Stage 2	Live market-state API and alert service for a narrow market universe.
Stage 3	Institutional strategy-suitability and risk-overlay module integrated with client workflows.
Stage 4	Client-specific allocation, execution or portfolio-control extensions subject to regulation and validation.

COMMERCIAL DISCIPLINE

The PhD should create reusable data assets, methods and evaluation evidence. It should not depend on the assumption that one proprietary strategy will generate persistent excess returns. The business case is stronger when the platform improves decision context, model governance and strategy selection across multiple clients and strategies.

13. Proposed Chapter Structure

Chapter	Title	Purpose
1	Introduction	Research context, problem, aim, questions and contribution.
2	Literature Review	Regime detection, market microstructure, ensemble learning, NLP, agents, uncertainty and governance.
3	Research Design and Data Architecture	Point-in-time data, feature lineage, taxonomy and experimental controls.
4	Multidimensional Market-State Detection	Development and comparison of specialist regime models.
5	Calibrated Specialist Evidence	Probability calibration, conditional reliability and common evidence schema.
6	Macro-Event and Narrative Intelligence	Structured textual information and incremental-value evaluation.
7	Multi-Agent Hypothesis Reconciliation	Consensus, disagreement, contradiction and competing state hypotheses.
8	Regime-Conditioned Strategy Selection	Strategy suitability, exposure and deterministic risk controls.
9	Integrated Evaluation and Commercial Prototype	Ablations, robustness, real-time prototype and business blueprint.

Chapter	Title	Purpose
10	Governance, Conclusions and Future Research	Auditability, limitations, findings and extensions.

14. Indicative Three-Year Research Plan

Period	Indicative activity
Year 1: foundations and market-state detection	Complete literature review; finalise taxonomy; acquire and govern data; build point-in-time platform; establish baselines; develop trend, volatility, liquidity and transition models; prepare first paper.
Year 2: evidence calibration and hypothesis reconciliation	Develop probability calibration and reliability framework; build macro-event agent; compare ensemble and consensus methods; develop live research dashboard; prepare second paper.
Year 3: strategy suitability, integration and commercial prototype	Evaluate regime-conditioned strategies and risk controls; test adaptive weighting; complete ablations and robustness analysis; develop institutional prototype and commercialisation chapter; submit thesis.

15. Scope, Limitations and Research Risks

15.1 Scope Controls

- Begin with one liquid market ecosystem rather than attempting universal cross-asset coverage.
- Use one-minute decision intervals initially; add tick-level components only where they are necessary for a defined research question.
- Treat the LLM as a structured event-information processor, not an autonomous trader.
- Evaluate market-state inference and strategy suitability before direct trade optimisation.
- Introduce RL only where transparent baselines have been established.
- Prioritise calibration, robustness and attribution over headline returns.

15.2 Principal Risks and Mitigations

Risk	Mitigation
Regimes are unstable or subjective	Use multiple operational definitions; report sensitivity; distinguish economic interpretation from statistical clustering.
Look-ahead and revision bias	Enforce available-at timestamps, release vintages, delayed execution and frozen historical information sets.
Data cost and licensing	Begin with a narrow market and scalable data tiers; document substitute datasets and limitations.
Agent complexity without value	Require ablation evidence and compare against simple ensembles and rules.
Uncalibrated confidence	Use formal probability calibration and reliability testing.
LLM hallucination or drift	Use schema constraints, retrieved evidence, source attribution, versioning and deterministic validation.
RL instability	Delay RL; restrict action spaces; use offline evaluation and hard risk constraints.

Risk	Mitigation
Backtest overfitting	Use walk-forward testing, multiple-testing controls, PBO, Deflated Sharpe Ratio and economic stress tests.
Latency undermines real-time use	Measure end-to-end latency and separate low-frequency narrative processing from fast numerical inference.
Commercial value depends on profitability	Evaluate decision usefulness, risk reduction, governance and workflow value independently of direct alpha.

16. Concluding Proposition

The revised proposal is not principally a claim that a collection of agents can predict market direction. It is a study of whether heterogeneous evidence can be transformed into a calibrated, multidimensional and auditable understanding of the current market environment.

Its central innovation is the separation of perception, evidence, hypothesis reconciliation, strategy suitability and risk control. Quantitative models describe specific aspects of the market. Specialist agents validate and standardise their outputs. A consensus layer compares competing explanations. Strategy agents determine whether particular approaches are appropriate. Deterministic controls retain final authority.

This positioning is academically more defensible than a broad autonomous-trading proposal and commercially more valuable than a single strategy. The resulting PhD could produce both publishable research and the technical blueprint for an Adaptive Market Intelligence company.

RECOMMENDED POSITIONING
Frame the doctorate as research into probabilistic market-state inference and regime-conditioned decision intelligence. Treat trading performance as one downstream evaluation, not as the sole proof of value.

Appendix A: Key Formal Definitions

Concept	Illustrative formulation or interpretation
Market-state vector	z_t = [trend_t, volatility_t, liquidity_t, driver_t, transition_t, cross_asset_t]
Specialist evidence	e_i,t = {state, probability, calibrated probability, reliability, horizon, freshness, support, contradiction}
Calibration condition	P(Y=1 | predicted probability=p) approximately equals p.
Competing hypotheses	H_k represent alternative explanations of the current market; posterior probabilities are updated from specialist evidence.
Belief state	b_t(s) = P(S_t=s | observations and prior information).
Transition probability	P(S_{t+1}=j | S_t=i) or an online probability that the current state is changing.
Strategy suitability	u_j,t = expected utility of strategy j conditional on the inferred state and risk constraints.
Risk-adjusted objective	Utility may penalise variance, drawdown, Expected Shortfall, turnover, transaction costs and liquidity risk.
Expert reliability	w_i,t may depend on historical calibration, similar-state performance, freshness and distribution-shift indicators.
No-action condition	Action is withheld where confidence, data quality, liquidity or distributional validity falls below defined thresholds.

Appendix B: Minimum Viable Research Prototype

Element	Initial implementation
Market universe	NQ futures, QQQ, NASDAQ breadth, volatility index, US Treasury yields and dollar proxy.
Decision interval	One minute.
Data stack	Python, Parquet or object storage, PostgreSQL or TimescaleDB, streaming layer where needed, MLflow and FastAPI.
Initial specialists	Trend, volatility, liquidity, structural change, macro-event and cross-asset.
Consensus baselines	Equal weighting, majority vote, Bayesian model averaging and supervised stacking.
Strategy families	Trend, mean reversion, long volatility, market neutral and cash preservation.
First user output	Live market state, confidence, disagreement, evidence, strategy suitability and risk recommendation.
First research milestone	Demonstrate that market-state information improves at least one of calibration, drawdown, tail risk or strategy selection out of sample.