From Probability to Coherence - A Structural Perspective on Market States

Introduction

Modern market analysis has largely been framed in terms of probability. From classical time-series models to advanced machine learning systems, the dominant paradigm seeks to estimate the likelihood of future outcomes based on historical data. Whether through regression, regime-switching models, or probabilistic classification, the objective remains consistent: infer expected returns from historical patterns.

This probabilistic lens has delivered enormous progress. It has formalised risk, quantified uncertainty, and enabled scalable decision systems across asset classes. Yet it rests—implicitly or explicitly—on a foundational assumption: that historical relationships retain sufficient stability to inform the future.

Markets, however, are not static systems. They evolve. Liquidity structures shift. Policy regimes change. Participant behaviour adapts. Technological transformations alter microstructure. In such environments, statistical relationships can degrade, invert, or fragment.

This paper introduces a complementary perspective: structural coherence.

Rather than asking: *"What is the probability of the next outcome given the past?"*,

Structural coherence asks a different question: "*Is the present configuration of market evidence internally consistent with a defined regime thesis?"*

This reframing shifts the objective:

- From prediction to validation
- From estimating likelihood to evaluating alignment
- From optimising expected return to assessing structural integrity

The distinction is subtle—but consequential.

The Probabilistic Paradigm

Over the past several decades, quantitative finance has been shaped by probabilistic reasoning. From early factor models to contemporary machine

learning systems, the central task has remained consistent: estimate future outcomes using patterns embedded in historical data.

In its simplest form, this involves modelling expected returns as a function of observed variables—such as momentum, volatility, liquidity, macro indicators, or structural factors. More advanced frameworks extend this logic into regime-switching models, hidden Markov processes, or probabilistic classification systems. Even when modelling market states rather than raw returns, the objective is typically to infer the probability of being in one regime versus another.

This probabilistic approach rests on an important but often understated premise: that the statistical relationships observed in the past retain enough continuity to inform the future. This does not imply markets are permanently stable. Most modern models account for volatility clustering, regime transitions, and time-varying parameters. Even adaptive systems ultimately rely on historical estimation, assuming that recent data retains sufficient informational value to guide near-term expectations.

In essence, probabilistic models ask: "*Given what has historically followed patterns like this, what is most likely to happen next?"*

This question is powerful. It has shaped portfolio construction, risk modelling, and algorithmic execution at scale. But it is not the only meaningful question one can ask.

A Founder's Shift: From Prediction to Structural Coherence

My background is not rooted in advanced statistical modelling. I did not approach markets by first studying probability theory, econometrics, or machine learning. Instead, my engagement with markets evolved through practical involvement in US equity markets—combining active trading experience with the use of multiple data platforms and early generative AI tools to develop structured insights. This eventually led to the iterative development of an automated insights framework using Generative AI.

As I worked on the framework, I noticed that many decisions were not truly about predicting price. They were about assessing whether a thesis about the market state made structural sense.

At various moments, the question was not: *“What is the probability of the next move?”*

It was:
“Is the current configuration of evidence internally consistent with this market state?”

and also: *“What if the primary objective is not to predict outcomes, but to validate whether a market state is structurally coherent?”*

That distinction became increasingly important. I began defining structured market-state combinations using price, technicals, fundamentals, earnings, news, insider sentiment, and contextual signals.

Over time, additional questions emerged:

- Does any other market state support this structure?
- Do contextual conditions reinforce or weaken it?
- Are there explicit invalidation criteria?

These questions gradually shaped a different design approach.

The shift was not from quantitative thinking to something else. It was from focusing primarily on likelihood to focusing on coherence.

The framework evolved toward layered alignment, where confidence builds progressively as structural components reinforce one another—signals, context, confirmation, and invalidation criteria forming a coherent whole.

This reframes decision-making from probabilistic inference to structural validation. In probabilistic systems, confidence emerges from historical frequency. In structural coherence systems, confidence emerges from alignment strength.

Markets as Evolving Systems: Non-Stationarity and Regime Transitions

Financial markets are not static environments. They evolve as capital flows shift, macroeconomic conditions change, regulatory frameworks adapt, and participants respond to new information and incentives.

Statistically, markets are often described as non-stationary systems. This means that the underlying relationships between variables—returns, volatility, correlations, liquidity dynamics—do not remain constant over time.

Parameters estimated from one period may not fully describe another.

- Momentum regimes strengthen and weaken.
- Liquidity dynamics evolve.
- Volatility compresses and expands in cycles.

What appears reliable under one configuration can behave differently under another.

Most quantitative frameworks account for this reality to some degree. They incorporate rolling windows, adaptive parameters, volatility clustering, or regime-switching mechanisms. Yet even these adaptive systems ultimately rely on historical estimation.

Non-stationarity becomes most visible during regime transitions.
These are moments when:

- Liquidity structures shift
- Policy regimes change direction
- Risk appetite resets
- Technological or structural shifts alter participation dynamics

During such transitions, previously stable statistical relationships can weaken or invert. Patterns that once carried predictive value may lose reliability—not because they were incorrectly modelled, but because the system itself has reconfigured.

A structural perspective approaches these moments differently. Rather than relying primarily on historical frequency, it evaluates whether the present configuration of signals, context, and confirmation remains internally coherent.

The question becomes less about whether a pattern has historically worked, and more about whether the structural conditions that would justify its continuation still exist.

This distinction shapes how decision systems respond when markets are not merely fluctuating within established dynamics—but evolving or redefining them.

Implications for Decision Design

If markets are evolving systems, and if regime validity can degrade as structural relationships shift, then the design objective of decision frameworks changes.

The central question becomes: “What should a decision system optimise for?” Traditional probabilistic systems optimise for expected value—maximising return relative to estimated risk based on historical data.

A structural coherence framework optimises for something different:

- Alignment of evidence
- Explicit thesis validation
- Clear invalidation criteria
- Gradual conviction scaling

This shifts the emphasis from prediction to architecture.

In practice, this means designing systems that:

- Define regimes explicitly rather than infer them implicitly
- Evaluate multiple dimensions of evidence—signal, structure, confirmation, and context
- Incorporate clear invalidation thresholds to manage regime breakdown
- Treat conviction as a function of alignment strength rather than solely historical frequency

Under this lens, the goal is not to forecast the precise magnitude of the next move. It is to determine whether the structural conditions supporting a thesis remain intact. Such systems prioritise interpretability, robustness during transitions, and clarity in risk governance. The difference is not merely methodological. It is philosophical. One approach seeks to maximise predictive accuracy. The other seeks to preserve structural integrity. Both have value. But they optimise for different objectives.

In the Context of Institutional Strategies

Institutional-level trading strategies are often categorised by how they extract alpha. These may include:

- Algorithmic trading
- Arbitrage
- Market making

- Long/short equity
- Factor-based systematic strategies

Each of these approaches focuses primarily on execution methodology, signal extraction, or capital deployment structure.
A structural coherence framework operates at a different layer.

It does not replace algorithmic execution, arbitrage logic, or long/short positioning. Instead, it functions as a regime-aware validation layer that can sit above or alongside these strategies.

In practical institutional settings, structural coherence can serve as:

- A gating mechanism for capital deployment
- A regime filter for systematic models
- A structural risk overlay
- A thesis validation framework prior to position sizing

For example, a long/short equity desk may generate signals based on valuation spreads and factor exposures. A structural coherence layer can evaluate whether broader market conditions reinforce or weaken the regime assumptions underlying those signals.

Similarly, a systematic momentum strategy may remain statistically valid in isolation. A structural overlay can assess whether liquidity, volatility, and contextual conditions align with continuation dynamics before full exposure is deployed.

In this sense, structural coherence is not an execution strategy. It is a decision architecture layer.

Its role is not to determine how trades are executed, but rather to determine whether the structural environment justifies participation in a defined regime.

Toward Integrated Frameworks

Market analysis does not require a single dominant lens. Different approaches serve different objectives.

Probabilistic models provide structured estimation of expected outcomes under defined assumptions. They remain foundational in portfolio construction, risk modelling, and systematic strategy design.

A structural coherence framework serves a different role. It formalises regime definitions, evaluates evidence alignment, and enforces explicit invalidation logic. Its primary contribution lies in decision architecture—how conviction is formed, how risk is governed, and how regime integrity is maintained.

In practice, these approaches need not be mutually exclusive. Statistical inference can inform expected return distributions. Structural coherence can govern whether a regime thesis remains valid.

Markets will continue to evolve. Frameworks designed with clarity of purpose—whether probabilistic, structural, or integrated—are better positioned to adapt alongside them.