



WHITE PAPER

Proof Over Performance

A practical governance framework for evaluating
AI in institutional investment

Michael King — Founder & CEO, QuantML | MBA, Trinity College Dublin

Proof Over Performance

A practical governance framework for evaluating AI in institutional investment

Michael King — Founder & CEO, QuantML | MBA, Trinity College Dublin

A QuantML white paper · July 2026

Artificial intelligence is entering investment decision-making faster than the frameworks used to govern it. This paper argues that the first question institutions should ask of any AI system is not how well it has performed, but whether its process can be trusted, reproduced and defended. It sets out five questions every investment committee should be able to answer before allowing an AI system near client capital.

The wrong first question

When a new investment capability arrives, the instinct is to ask a single question: how well does it perform? With artificial intelligence, that instinct is misplaced. A track record — particularly one produced by a model — can look extraordinary and still tell you almost nothing about whether the process behind it can be trusted. A model can generate an impressive set of results and remain impossible to reproduce, impossible to audit and impossible to explain. For an institution, that combination is not an asset. It is a liability waiting to surface, at precisely the moment a client or regulator asks how a decision was made.

The more useful first question is different: can you prove how it works, and would that explanation withstand scrutiny? Performance is an output. Proof is a property of the process. The next frontier in institutional AI is not better performance — it is stronger proof.

A governance problem, not a technology problem

Investment committees already know how to evaluate a manager. They examine the process, not just the outcome; they ask whether results are repeatable; they probe the risk controls; they establish who is accountable. Artificial intelligence should be held to that same standard — no lower because the technology is novel, and no higher because it is unfamiliar.

Seen this way, the problem with a black box is not that it uses machine learning. It is that a black box cannot answer the questions a committee would put to any manager. The remedy is not to fear the technology, nor to take it on faith, but to extend the due diligence that already exists. The five questions that follow are that extension.

Five questions every investment committee should be able to answer

The following framework is deliberately provider-agnostic. It is a lens for evaluating any AI investment platform — including ours — and it is written so that every member of a committee, technical or not, can use it. Each pillar comes with the questions we believe a committee should ask, and be given clear answers to.

1. Transparency

Transparency is the ability to see how a decision was reached — the logic, the inputs and the reasoning behind an output, expressed in investment terms rather than technical ones. Opacity is usually a choice, not a technical necessity.

Questions to ask:

- Can the provider explain, in plain investment language, why a particular signal or decision was produced?
- Are the inputs, and the sources of those inputs, disclosed?
- Could the model's behaviour be described to a client or a board without resorting to hand-waving?

2. Reproducibility

Reproducibility is the ability to obtain the same output from the same inputs, every time. A result that cannot be reproduced is an anecdote, not evidence — and anecdotes have no place in an institutional process.

Questions to ask:

- If the same process were re-run on the same data, would it produce the same result?
- Is every published figure traceable to an underlying record rather than a summary?
- Could an independent party reconstruct the result from first principles?

3. Verifiability

Verifiability is the ability of a third party to check a claim independently. Trust should never require faith. The phrase that matters here is a capability, not a slogan: results should be independently verifiable and reproducible.

Questions to ask:

- Can a claim be checked without taking the provider's word for it?
- Is there a timestamped record that predates the outcome being claimed?
- Is it clear exactly what is being verified — the signal, its settlement, or the performance derived from it?

4. Auditability

Auditability is a complete, tamper-evident trail from data to decision to outcome. It is what allows a committee, a client or a regulator to follow the chain end to end, without gaps and without having to trust the narrator.

Questions to ask:

- Is there a full audit trail linking every output back to its origin?
- Does each figure trace to a settled, canonical record, with no restatements after the fact?
- Could the whole chain be handed to an external auditor tomorrow?

5. Governance

Governance is the human and procedural scaffolding around the model: version control, change management, risk controls and clear lines of accountability. A model without governance is an experiment; a model with governance is an institutional process.

Questions to ask:

- How are changes to the model governed, recorded and reviewed?
- What risk controls apply, and do they adapt as conditions change?
- Who, precisely, is accountable when the model is wrong?

The five-question test

Distilled, the framework is a single test a committee can apply to any AI provider in an afternoon. If the provider cannot give clear, evidenced answers to these five, the performance record is beside the point:

1. Transparency — Can you show me how a decision was reached?
2. Reproducibility — Can you produce the same result from the same inputs?
3. Verifiability — Can I check your claims without taking them on trust?
4. Auditability — Can you trace every figure to a settled record?
5. Governance — Can you tell me who is accountable and what controls apply?

Why this matters now

Adoption of AI across investment management is accelerating, and client scrutiny and regulation are following close behind. As they arrive, the firms that earn institutional trust will not be those with the boldest performance claims. They will be those that can demonstrate proof — a process that can be seen into, reproduced, verified, audited and governed.

Being able to answer these five questions is fast becoming the price of entry rather than a point of differentiation. The standard is coming either way. The only real choice is whether to adopt it before clients and regulators require it, or after.

At QuantML we built our platform around these five principles from the first line of code, in the conviction that evidence-based, independently reproducible investment research is what the coming decade of institutional AI will demand. But the framework is not ours to keep. Use it to evaluate anyone who brings AI near your clients' capital — ourselves included. That is the point.

About the author

Michael King is Founder & CEO of QuantML and holds an MBA from Trinity College Dublin. He has more than 25 years' experience across technology, risk and operational transformation in global financial institutions. QuantML is an Irish institutional AI investment intelligence platform built around transparency, auditability and reproducible investment research. Michael is a recipient of a Fingal Local Enterprise Office Feasibility Grant supporting the development of QuantML (March 2026).

mking@quantml.ai · quantml.ai · © 2026 QuantML Limited