



Can Your AI Explain Itself? What Black Box AI Means for Financial Organizations

By Michael Berman

Artificial intelligence is making decisions inside your financial organization right now. It's screening transactions, flagging risk, and surfacing exceptions. In many cases, it's doing it faster than any human review process could.

But let's say an examiner asks how your fraud detection system reached a particular conclusion: what it weighed, what guidance it applied, and whether the decision can be defended. You pull up the output. The answer is there, but there's no source, no reasoning, and no record of how the model arrived at its conclusion.

That's the black box AI problem. For the 22% of financial organizations that have adopted AI in compliance work, it's not just hypothetical — it's a major barrier. Addressing it starts with understanding what black box AI is and ensuring the technology your organization relies on can meet the explainability standard examiners expect.

When “The AI Told Us” Isn't Enough

A black box AI system produces outputs without explaining how it got there. For general productivity tasks, that trade-off might be acceptable, but for compliance decisions that carry legal, operational, and reputational weight, not knowing how a conclusion was reached is its own form of risk.

The outputs that flow from compliance decisions shape how your organization manages risk, prepares for examinations, trains staff, and demonstrates adherence to the regulations governing your existence. When those outputs can't be explained, you don't just have an AI problem; you have a documentation problem, a governance problem, and potentially an examination problem all at once.

General-purpose large language models compound this dilemma. Trained on broad public data and optimized for confident, fluent responses, LLMs don't know your charter type, supervisory history, or regulatory environment. They can't distinguish current guidance from guidance that was superseded two years ago. Worse, they have no mechanism for recognizing the limits of what they know, so they'll produce authoritative-sounding answers to questions where the correct response is nuanced or jurisdiction-specific.



The danger isn't that LLMs are always wrong. It's that they're wrong in ways that are hard to catch without deep subject matter expertise, and they leave no traceable record of how the answer was produced.

What Examiners Expect

AI regulatory guidance has been building for years across the OCC, FDIC, CFPB, Federal Reserve, NCUA, Fannie Mae, and Freddie Mac. No unified standard exists yet, but the underlying expectations come down to three things: explainability, accountability, and auditability.

- **Explainability** requires being able to articulate what inputs informed an AI output, what sources were consulted, and why the conclusion holds up.
- **Accountability** means someone at your organization owns the output. Management remains responsible for risk management processes regardless of whether they are technology-assisted.
- **Auditability** gives examiners a chain to follow from question to answer, from input to output, from AI-generated material to the human decision it influenced.

How to Rethink Your AI Approach

Meeting those three expectations requires AI that's built that way from the start. That's what glass box AI is: a system designed to show its work. Sources surface with every response, and the reasoning is transparent enough that a trained compliance professional can evaluate the output, verify its accuracy, and document the basis for any decision that follows.

Before deploying any AI tool in compliance or risk functions, your organization must be able to answer a few critical questions:

- Can the system identify the specific regulatory sources it drew upon?
- Are its responses verifiable against those sources?
- Can you log and retain inputs and outputs in a retrievable format?
- Is it calibrated to your regulatory environment, or generating generic answers from undifferentiated public data?

If any of those can't be answered, you have risk exposure.



The Documentation Standard

Documentation in an AI-assisted compliance environment goes beyond saving a copy of the output. The standard should capture what was asked, what documents were submitted, the complete response, the specific regulatory texts cited, and who reviewed it, including what they evaluated it against and whether it was modified before use.

For regulatory research, retain the original query, the full response, the sources cited, the reviewer's name, and any changes made before the output was relied on. Store it somewhere retrievable for an exam, not buried in a chat thread or a personal folder. When AI plays a role in exam preparation or a control assessment, document the scope of the review, where AI output was incorporated, and how it was validated before use.

When an examiner starts asking questions, documentation is what separates a program that holds up from one that doesn't.

When the Examiner Asks, You Need an Answer

The organizations best positioned for AI-related exam scrutiny won't necessarily be the ones that moved slowest. They'll be the ones that were deliberate about which AI they deployed, how they used it, and what controls surrounded it.

When an examiner asks how your organization reached a specific conclusion, the question isn't whether you used AI to get there. It's whether you can open the box and show them exactly what happened.

About Ncontracts

Ncontracts provides integrated risk management, compliance, and third-party risk management solutions to over 5,500 organizations worldwide, including 4,500 U.S. financial institutions, mortgage companies, and fintechs. The flagship Ncontracts IRM suite combines AI-powered software with expert services, helping financial institutions streamline risk, compliance, and vendor management through an intuitive, cloud-based platform. Ncontracts' Venminder solution is trusted by enterprise financial companies and other large organizations to strategically manage third-party risk across the entire vendor lifecycle.

Visit www.ncontracts.com or follow the company on [LinkedIn](#) and [X](#) for more information.