

When Your AI Can't Show Its Work, You Own That Lending Risk

By Rafael DeLeon

Lenders are using artificial intelligence (AI) well beyond back-office compliance work: underwriting models, fraud detection, pricing engines, and loan origination platforms. Every time you deny credit or offer an applicant less favorable terms, there's a reason behind it. A risk score or a flag might be part of that answer, but it isn't the whole answer. It must be something the applicant can understand and something you can defend.

That's the black box problem: when an AI system makes a decision with no reasoning behind it. A glass box, on the other hand, shows its work so that someone qualified can evaluate and approve the output.

Why Black Box AI is Lending Liability


Under the Equal Credit Opportunity Act (ECOA) and Regulation B, you must give applicants specific, intelligible reasons for denying credit or offering less favorable terms. If your AI can't trace a decision back to a clear rationale, "the model said so" isn't a defense; it's a failure to explain that compounds once regulators get involved.

HMDA adds another layer of exposure for mortgage lenders. Regulators can see denial and pricing patterns across your portfolio, and simply pointing to the model won't answer questions that are raised.

Fair Lending Risk Didn't Go Away

In April 2026, the CFPB removed disparate impact from ECOA enforcement. You may read that as reduced pressure, but the Fair Housing Act still carries disparate impact liability, and the Supreme Court precedent underlying it remains intact.

State enforcement fills the gap federal posture leaves open. Massachusetts settled a case last year for \$2.5 million tied to AI-driven underwriting that produced disparate outcomes, and New Jersey finalized rules naming AI and automated decision-making as sources of disparate impact risk.



If you operate across state lines, this exposure holds regardless of where federal guidance lands, and it isn't limited to underwriting. If similarly situated borrowers got different rates because AI shaped pricing (or an applicant was steered toward a particular product), you need a documented, legitimate rationale behind it.

The Standards Are Being Set

Model risk governance has evolved over the past year. A major change came on April 17, 2026, when the Federal Reserve, OCC, and FDIC replaced the [model risk management guidance](#) that had been in place since 2011. It's aimed primarily at banking organizations with over \$30 billion in total assets, but that doesn't mean smaller FIs are off the hook.


A third-party model is still *your risk*, regardless of your size. That matters most for smaller banks, often the heaviest users of vendor-built tools such as allowance for credit losses (ACL), asset/liability management (ALM), fraud detection, and anti-money laundering (AML) monitoring. If you rely on a vendor model, you need to understand how it works, validate it, document it, and monitor it over time. Examiners expect a clear line of sight from how the model is used to who's responsible for it.

Government-sponsored enterprises (GSEs) have also raised expectations. [Freddie Mac's requirements for seller/servicers](#) took effect March 3, 2026, and [Fannie Mae's framework](#), issued in April, took effect August 6, 2026.

What a Documentation Standard Should Cover

Saving the output from an AI tool isn't enough. A defensible record needs the original question or prompt, anything submitted as input, the complete response, the specific sources or factors relied on, and a note on who reviewed it and what changed before it was used. That record must hold up throughout the model's lifecycle, not just at deployment.

This doesn't mean re-underwriting every loan by hand, but it does require a human in the loop. Risk analysts, compliance leaders, and underwriters must set up model parameters, review outputs, and flag when a model's behavior warrants a second look.



What Glass Box AI Delivers

Not all AI is equally explainable. Traditional statistical models and many machine learning approaches map back to identifiable factors, weights, and thresholds. More complex models may support explainability, but only if you understand how inputs drive outputs and can translate that logic into reason codes an applicant or examiner can follow.

Glass box AI is what that looks like in practice: an auditable, traceable system that leaves a record of what went in, what came out, and why in a form a human can follow.

Before relying on any AI touching a credit decision, you must be able to answer who owns the model and its data, how it's tested and validated, whether the output can be explained in plain terms that satisfy Regulation B, and who's monitoring it and for what.

Vendor contracts need to cover the same ground. If an agreement doesn't clearly address AI use, explainability, and information-sharing obligations, start there.

The Bottom Line

Explainability isn't about opening the black box for its own sake. When a credit decision affects a consumer, you need to articulate why the outcome occurred and stand behind it. As AI takes on a larger role across the lending lifecycle, that's not an aspiration. It's the baseline.

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.