



From automated underwriting to artificial intelligence

How fair lending risk oversight must evolve

BY REBECCA ESCARIO, CAMS AND JONATHON NEIL

Machine learning has been a component of lending for decades, powering everything from credit scores to automated decision engines. The financial industry is not new to algorithmic models; however, the landscape is shifting dramatically. Today's artificial intelligence (AI) brings more complex models and a new set of fair lending risks that demand thoughtful oversight.

One of the most pressing concerns in this new landscape is model explainability. Unlike traditional rules-based systems, many AI and machine learning models function as "black boxes," challenging lenders and regulators to clearly understand how input data is being used to generate decisions. This opacity can hinder efforts to detect or correct discriminatory practices.

Another significant new risk comes from the potential use of proxy variables. Even when sensitive variables such as race or gender are explicitly excluded, other data points can serve as proxies that unintentionally introduce or perpetuate bias. Dynamic updating of systems also presents a unique challenge. AI models are increasingly capable of learning and evolving over time, updating decision criteria as they are exposed to new data. While AI systems can

adapt and learn, improving both efficiency and accuracy over time, banks must remain vigilant in monitoring these evolving patterns. This ongoing oversight ensures that algorithmic changes do not inadvertently create disparate impacts on protected consumer groups. Ongoing oversight, robust documentation, and continuous testing are necessary to ensure these evolving models remain fair, transparent, and compliant.

The risks mentioned above are also rapidly evolving in the context of bank-fintech partnerships. These collaborations are transforming the financial services and regulatory environments, creating new opportunities and challenges; however, they also introduce substantial fair lending risks that require diligent management.

The shift from traditional automated underwriting systems (AUS) to advanced AI represents an evolution, not a revolution — a balance of continuity and change. AUS and other rules-based engines were the first wave of algorithmic underwriting. Modern AI is simply the next generation, sharing the same fundamental goals, but operating with far greater complexity and opacity while potentially introducing unintended bias.

The core challenge for compliance professionals is not that AI is an entirely new concept. Instead, the challenge lies in adapting oversight methods. As relatively transparent rules-based systems give way to sophisticated AI-driven decision making, a bank's approach to managing fair lending risk must evolve alongside the technology. The core principles of fairness remain, but the tools and techniques we use to uphold them must be updated for this new era of lending.

Understanding the sources of AI fair lending risk

To effectively manage fair lending risks associated with AI, banks must understand their origins. Risks can emerge at every stage of the AI lifecycle: development, implementation, and ongoing governance. Each phase presents unique challenges that require dedicated and thoughtful oversight. (See Figure 1: AI Lifecycle Phases.)

- **During model development:** The potential for bias is significant as a model is constructed. The data used to train AI models can propagate, or even amplify, historical biases, such as those resulting from past redlining practices. If training data is skewed, a model's decisions will be as well. The selection of data features also poses a

risk; for example, using alternative data such as rental or utility payment history can expand credit access, but it may also introduce new proxy variables for protected characteristics if not carefully validated. Further, the use of non-traditional data points, such as alternative credit history metrics, occupation, education levels, social media activity, and online browsing and purchase history, raises serious ethical questions about privacy and consent, as these data sources may not have a clear connection to creditworthiness.

- **During model implementation:** Once a model is developed, its implementation creates another layer of risk. Decisions on how to integrate the AI's output, such as setting specific cutoffs or loan-level price adjustments, can introduce bias even if the underlying model is sound. While some of these integration variables may elevate risks in ways similar to traditional underwriting, integrating AI model output may warrant setting new thresholds and cut-offs and therefore introduce new risk into the credit decisioning process. A common technical issue is overfitting, where a model performs well on its training data but fails to generalize to new applicants, leading to poor and potentially discriminatory outcomes.


- **During model governance:** Fair lending risk may enter during oversight of the model process also. AI models are not static; their performance can drift over time as economic conditions and applicant pools change, potentially creating new disparities that go unnoticed without active monitoring. A lack of transparency, or "explainability," makes it difficult to understand why a model made a certain decision, complicating compliance insufficient board-level governance. Increasingly, financial institutions are partnering with fintech firms to develop and deploy AI models, which raises the stakes for

strong governance. Success in these partnerships requires robust oversight, both to satisfy regulatory expectations and to assure fintech partners that their innovations will not expose them to unnecessary compliance risk.

AI-driven targeted marketing and applicant engagement tools also bring a new dimension of risk to AI model governance. For example, machine learning algorithms that determine which consumer groups receive online credit card advertisements or loan offers might inadvertently exclude certain demographic groups, particularly if they rely on data tied to zip codes, browsing histories, or social media. AI-powered lead generation systems may steer marketing efforts toward or away from applicants based on learned patterns in historical applicant pools, reinforcing gaps in access. Even "digital loan officers" and chatbots designed to guide applicants through the application process can subtly shape user experience by providing more information, tailored encouragement, or different service levels to certain users. These practices may influence not only who enters the application pipeline, but also how welcomed, informed, or encouraged prospective borrowers feel, making outcome monitoring and careful oversight of AI-driven marketing and support platforms essential for compliance.

The shift from traditional automated underwriting systems to advanced AI is an evolution, not a revolution — a balance of continuity and change.

FIGURE 1: AI Lifecycle Phases



Key considerations for monitoring and compliance

As AI reshapes lending, compliance, and monitoring, frameworks must evolve to address new forms of risk. A primary challenge is the classic tradeoff between model accuracy and transparency. Complex, "black box" models may deliver superior predictive power but can obscure the exact factors driving a decision, making it difficult to prove that a decision was not directly or indirectly based on prohibited characteristics. This obliqueness complicates a bank's responsibility and requirement to provide accurate adverse action notices, a cornerstone of fair lending compliance. If a bank cannot explain why an applicant was denied, it cannot meet compliance obligations.

As mentioned above, the risk of using proxy variables should also be a key consideration for banks when implementing AI solutions.

AI models can inadvertently identify and use seemingly neutral data points – like a consumer’s brand of phone or their grocery shopping habits – that correlate strongly with protected classes. This practice has the potential to lead to disparate impact, where a neutral policy disproportionately affects a protected consumer group, without evident discriminatory intent.

The current regulatory environment is also a moving target, and the broader enforcement landscape is shifting. As federal agencies adjust their enforcement priorities, several state-level regulators are stepping in to fill the perceived gap, launching their own investigations and legal actions related to fair lending violations. Consumer protection advocacy groups at the national, state, and local levels are also increasing focus on fair lending. With the perception that federal oversight may be waning, these organizations are ramping up their own monitoring, research, and public awareness campaigns, adding another layer of scrutiny that banks must consider in their compliance and risk management strategies.

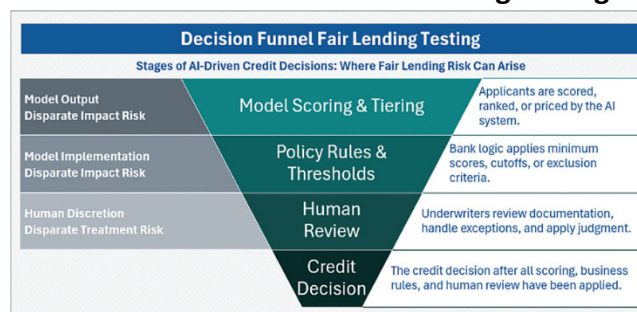
To remain compliant when implementing AI solutions, bank teams must maintain a robust and effective governance framework. Such a framework should include establishing cross-functional oversight, monitoring of marketing and initial applicant engagement, maintaining documentation of model development and validation, and conducting regular audits of both the AI models and their outcomes.

Even when sensitive variables are excluded, proxy data can quietly reintroduce bias — making vigilant validation and testing non-negotiable.

Practical steps for managing AI fair lending risk

The evolution of lending technology to AI-enabled solutions has not changed the basic fair lending compliance mandate: banks must be able to demonstrate that credit decisions are made fairly and without unlawful discrimination. Lenders should embed risk mitigation steps into every phase of model development, implementation, and governance.

FIGURE 2: Decision funnel fair lending testing



▪ **During model development:** The foundation of fair lending risk management begins before an AI model is put into production. Banks should test models throughout the development process for fairness by comparing alternative model forms and documenting how input variables are selected. Regulators expect banks to consider whether a less discriminatory alternative (LDA) could achieve comparable business objectives with lower disparate impact. Thoughtful evaluation of candidate models and transparent documentation of the trade-offs made is critical during this phase. Importantly, this evaluation should take place before the bank commits to a particular model family (logistic regression, random forests, etc.). Once development is underway, switching to a different model type can be prohibitively costly. Comparing model

families upfront helps to deter fair lending risk from entering into the final model while also strengthening LDA documentation.

The foundation of fair lending risk management begins well before an AI model is deployed. Banks should rigorously test models during development for fairness, comparing alternative model forms and thoroughly documenting the selection of input variables. While federal regulators' expectations around less discriminatory alternatives (LDA) have shifted — particularly with some agencies deprioritizing disparate impact cases — banks should remain vigilant. States, private litigants, and consumer groups may still pursue actions, especially under the Fair Housing Act, where disparate impact remains valid per the Supreme Court.

Regulators have historically expected banks to evaluate whether an LDA could achieve similar business objectives with reduced disparate impact. Conducting a thoughtful evaluation of candidate models and transparently documenting trade-offs remains critical. This process should occur before committing to a specific model family (e.g., logistic regression, random forests), as switching models mid-development can be prohibitively expensive. By comparing model families early, banks can mitigate fair lending risks in the final model and strengthen LDA documentation, ensuring compliance and reducing exposure to potential challenges.

Model development may include collaborating with professionals with limited compliance knowledge (data scientists, etc.); therefore, implementing fairness training and establishing guardrails for development decisions such as sample selection can reduce the likelihood that biases rooted in historical data carry forward

into new models. When working with third-party vendors, appropriate documentation is critical. Banks need to ensure they understand and can explain how a model was built and its key drivers, even if the underlying code and mechanics are proprietary.

▪ **During implementation:** A well-developed model can still create fair lending risk if it is poorly implemented. How AI model outputs — such as scores, automated decisions, or rates — are applied through established thresholds, exclusion rules, or overlays can impact disparities, often as much as or more than the raw model outputs themselves. Banks should test how these configuration decisions perform across different applicant groups and how they interact with manual overrides or exceptions that create unintentional patterns. Benchmarking against challenger models — explainable, policy-based factors in underwriting and pricing decisions — can help identify whether AI-enabled processes are changing, or increasing, fair lending risk.

▪ **During governance and monitoring:** Fair lending oversight continues after a model enters the production phase. AI models can drift over time as applicant characteristics, economic conditions, and market trends change. Layered statistical testing, including disparity analysis and regression models, at distinct stages of the decision process as shown in Figure 2, should be used to monitor whether disparities are emerging and at what point in the process they arise. Periodic audits of adverse action reasons should also be conducted to ensure that denial reasons are consistent with the actual drivers of the decision. Ongoing monitoring is particularly important for banks that deploy models developed by fintech partners. Vendor documentation may be incomplete and leave too

many unanswered questions about how the model operates and treats protected groups in practice. For these partnerships, banks should conduct independent outcome testing and establish vendor oversight expectations in contracts to ensure fair lending risk is identified and addressed, even when an underlying model remains proprietary. (See Figure 2: Decision funnel fair lending testing)

Fair lending cannot be an afterthought — it must be built into development, checked in implementation, and monitored continuously.

Banks should also evaluate how AI is used in marketing, lead generation, and applicant support tools, because targeted outreach or digital interactions can create redlining and steering risk. Vendor reviews are essential because third-party models do not release an institution from its fair lending responsibilities. Achieving strong governance in this evolving landscape includes alignment with ethical AI frameworks such as the National Institute of Standards and Technology AI Risk Management Framework to demonstrate the importance of transparency, accountability, and security. In practice, this also means ensuring robust cybersecurity measures to protect sensitive applicant data that AI models typically use.

The core principle is clear: fair lending cannot be an afterthought when deploying AI-enabled lending. Instead, it should be built into the development process, carefully considered during implementation, and monitored continuously.

How regulators may leverage AI themselves

While financial institutions are experimenting with or deploying AI-enabled solutions, federal agencies, including prudential regulators, have been directed to expand their use of AI as well. A new generation of economists and data scientists fluent in machine learning advances are increasingly joining supervisory agencies, bringing with them techniques that expand on legacy approaches.

According to a May 2025 Government Accountability Office report (Artificial Intelligence: Use and Oversight in Financial Services), financial regulators are already using AI in limited but expanding ways. Agencies are training staff, forming AI working groups, and collaborating with domestic and international institutions to build expertise. Current and planned applications include extracting and analyzing large volumes of examination documents, identifying risks and anomalies in supervisory data (Call Reports, lending disclosure filings, etc.), and flagging potential legal violations or reporting errors. While regulators indicate they are not relying on AI-made decisions, they are increasingly combining AI outputs with traditional supervisory information to improve efficiency and target exam priorities.

For fair lending, the implications are significant. If regulators begin applying machine learning and AI to analyzing lending data, they will be able to move beyond traditional disparity and regression-based approaches — which rely on relatively simple assumptions about the relationship between predictors and outcomes — to methods of pattern recognition that can flexibly capture non-linear statistical

relationships. These techniques may also help regulators cross-reference complaint data, Call Report data, market demographics, and other lending data, providing for more accurate outlier detection to identify potential fair lending risks with greater granularity.

Banks can prepare for these supervisory developments by anticipating how their Home Mortgage Disclosure Act, Community Reinvestment Act, and complaint data may appear when evaluated with advanced techniques. Preparation should include testing for anomalies or patterns across different segments (product, geography, etc.), benchmarking against similarly situated peers, and ensuring disparities are documented and explained in clear business terms.

Evolving oversight for an evolving world

AI represents an evolution in credit decisioning, presenting both opportunities and new fair lending risk. Traditional compliance tools must also evolve by embedding fairness testing in model development, implementation, and monitoring. Regulators are also developing their own AI capabilities that may soon change how examinations are conducted. Institutions that prepare now will be well-positioned to reduce regulatory risk, build trust with stakeholders, and succeed through fintech partnerships where strong governance is essential.

ABOUT THE AUTHORS

REBECCA ESCARIO, CAMS is a managing director at CrossCheck Compliance LLC and a compliance and risk management executive with over 25 years of experience with the financial and professional services industries. Previously Rebecca was the chief operating officer at a regulatory compliance consulting firm located in Fairfax, Virginia. Rebecca has consulted with banks and other financial institutions of all sizes on their fair lending, Home Mortgage Disclosure Act, Community Reinvestment Act, internal audit, and compliance needs. Rebecca can be reached at rescario@crosscheckcompliance.com.

JONATHON NEIL is a fair lending professional at CrossCheck Compliance LLC with over 12 years of consulting experience supporting financial institutions in the areas of Fair Lending and Community Reinvestment Act (CRA) compliance. His expertise includes data management, structured and unstructured data, mining and visualization, statistical modeling, red lining analyses, fair lending regression analyses, CRA performance assessments and community credit needs assessments. Jonathon can be reached at jneil@crosscheckcompliance.com.