

BRODY | GAPP LLP

Clarity in Compliance. Confidence in Litigation.

Fair Lending AI Risk Map for Mortgage

Where AI Creates Exposure — and What Controls You Need

Prepared by James W. Brody, Esq. | Founding & Managing Partner

James@BrodyGapp.com | 415.246.3995

AI does not need to be malicious to be discriminatory.

Proxy risk, data bias, and model opacity can recreate the effects of intentional discrimination at massive scale — in tools that were designed in good faith and selected in good faith. Freddie Mac Guide § 1302.8 imposes a binding, enterprise-wide obligation on every Seller/Servicer — including IMBs, depositories, credit unions, mortgage brokers, and fintechs — to govern, monitor, and audit every AI/ML tool used in connection with Freddie Mac loans, including vendor-embedded systems. Freddie Mac has not defined 'AI' or 'ML' in the Guide — the definitional ambiguity itself is a governance risk.

THE THREE SOURCES OF AI FAIR LENDING RISK

PROXY RISK

PROBLEM: A model feature that appears race-neutral — ZIP code, education, transaction patterns, credit utilization in LMI areas — correlates so strongly with race that using it produces the same discriminatory effect as using race directly. Prior redlined geographies are embedded in historical mortgage training data. The model learns those patterns without being told to look for race.

CONTROL: Pre-deployment proxy correlation analysis for every model feature. Remove or control for high-proxy features where no legitimate business justification exists. Document the business necessity analysis for any retained high-proxy feature.

OPAQUE MODELS (Black-Box Risk)

PROBLEM: Black-box ML models produce decisions without interpretable logic. ECOA § 701 and Reg B § 1002.9 require specific, actual reasons for adverse action. Note: CFPB Circular 2022-03, which had made this explicit, was withdrawn May 2025 — but the statutory obligations remain fully operative. A model that cannot produce ECOA-compliant reasons cannot legally be deployed for credit decisions, regardless of its predictive accuracy.

CONTROL: Model interpretability as a pre-deployment go/no-go condition. Map model feature importance directly to adverse action reason codes. Audit the mapping quarterly. Deploy only models from which you can extract specific, accurate adverse action reasons.

**DATA
BIAS**

PROBLEM: Historical mortgage data reflects decades of documented inequity — redlining, discriminatory appraisals, exclusionary lending practices. AI trained on this data learns those patterns and perpetuates them at scale, with no human ever deciding to discriminate. The bias is architectural, not intentional.

CONTROL: Pre-deployment fair lending analysis using BISG-estimated protected class where HMDA data is unavailable. Ongoing monitoring for approval and pricing disparities. Annual revalidation at minimum — earlier if market conditions shift significantly or model is retrained.

USE-CASE RISK MAP: WHERE TO LOOK AND WHAT TO WATCH

USE CASE	RISK TYPE	HOW IT SHOWS UP	CONTROLS
Income / Asset Verification AI	Proxy Risk Data Bias	Models trained on W-2 income data score gig workers and self-employed borrowers (disproportionately minority) lower. The model does not know it is disadvantaging non-traditional income — it learned that W-2 income is safer from past data.	Test approval rates by income type x demographics pre-deployment. Monitor gig vs. W-2 approval disparities on an ongoing basis.
AUS Overlays (AI credit policy)	Proxy Risk Opacity · Drift	An AI overlay applied on top of a DU or LP approval can create disparate impact without AUS accountability. Every overlay is a fair lending test you may not have run yet.	Fair lending analysis of overlay effect by race/ethnicity before deployment. Track overlay application rates vs. HMDA data. Document business justification for every overlay criterion.
Risk-Based Pricing Models	Proxy Risk Disparate Impact	AI pricing incorporating census tract, LMI credit utilization, or similar features produces pricing disparities that mirror historic redlining. HMDA data gives regulators and plaintiffs the data to find it.	Pricing disparity regression analysis using BISG-estimated race pre-deployment. Automated pricing disparity dashboard monitoring by demographic group on an ongoing basis.
Fraud Detection AI	Disparate Impact Opacity	Fraud models that flag thin-file or new-to-credit borrowers at higher rates disproportionately affect first-time minority buyers. A disparate false positive rate is a fair lending problem — one that few lenders are currently monitoring.	False positive rate testing by demographic group pre-deployment. Monitor referral-to-manual-review rates by race/ethnicity on an ongoing basis.
Document Classification AI (OCR / NLP)	Accuracy Bias	OCR and NLP models trained on standard documents may have higher error rates on non-English documents or non-standard income formats — creating disparate processing times and downstream outcomes.	Test extraction accuracy across document types and languages before deployment. Monitor error rates and exception handling by borrower demographics.

USE CASE	RISK TYPE	HOW IT SHOWS UP	CONTROLS
Marketing Segmentation AI	Digital Redlining	Marketing AI optimizing for conversion suppresses outreach to LMI or minority census tracts because they convert at lower rates. That is an FHA violation even if no discrimination was intended.	Geographic and demographic reach analysis before launch. Real-time redlining heatmap monitoring of marketing reach. Geographic exclusions require documented business justification.
Loss Mitigation Outreach AI	Proxy Risk Disparate Impact	AI that ranks borrowers for proactive outreach by default probability may systematically deprioritize minority borrowers for loss mitigation contact — a RESPA and FHA risk that is rarely tested.	Test outreach prioritization model for disparate rates by race/census tract pre-deployment. Monitor contact rates and loss mitigation outcomes by demographics on an ongoing basis.
Call Center AI Assistants	UDAAP Accuracy Risk	AI assistants providing incorrect information about forbearance options, modification terms, or payoff amounts — particularly to distressed borrowers — is a UDAAP problem regardless of intent.	Test accuracy across borrower scenario types before deployment. Human QC sampling of AI-assisted calls. Complaint monitoring with NLP categorization accuracy audit.
NLP Complaint Categorization	UDAAP Data Bias	An NLP model that systematically miscategorizes fair lending complaints as routine customer service issues suppresses regulatory complaint reporting and creates a compliance blind spot inside your own second line.	Test categorization accuracy specifically for fair lending complaint language. Periodic human audit of complaint categorization sample. Escalation protocols for any identified miscategorization pattern.
HMDA Anomaly Detection AI	Second-Line Integrity	If the AI tool used to monitor fair lending compliance itself has gaps or biases, your entire compliance monitoring function is compromised. The AI watching your AI needs to be validated too.	Independently validate any AI-based compliance monitoring tool. Cross-check AI-identified findings against manual analysis sample on a periodic basis.

EMERGING ISSUES TO WATCH IN 2026

- Alternative data in underwriting (rent history, utilities, cash-flow analysis) — fair lending implications largely untested at scale; regulators have not provided clear guidance;
- Generative AI in borrower communications — UDAAP risk if AI-generated content is inaccurate, misleading, or creates false urgency about loan terms;
- Shadow AI — employees pasting borrower NPI/PII into public generative AI tools (ChatGPT, Gemini, etc.) — data security risk and potential Freddie Mac § 1302.2 cybersecurity policy violation;

- State AI regulation patchwork — Colorado AI Act (June 30, 2026), Texas TRAIGA (January 1, 2026), and additional state action expected; federal preemption efforts underway but unresolved;
- Interagency AVM Final Rule (effective October 1, 2025) — requires nondiscrimination standards for automated valuation models used by mortgage originators and secondary market issuers; backed by Dodd-Frank § 1473(q) — more durable than guidance documents;
- Digital redlining litigation — the DOJ housing ads case established that algorithmic marketing discrimination is actionable under the Fair Housing Act without discriminatory intent; and
- Fannie Mae asymmetry — Fannie Mae has not issued an AI governance mandate; dual-eligible lenders should build to the Freddie Mac standard as the enterprise baseline

James W. Brody, Esq.

Founding & Managing Partner
BRODY | GAPP LLP

415.246.3995 · James@BrodyGapp.com · www.BrodyGapp.com

*General information only — not jurisdiction-specific legal advice. This material does not establish an attorney-client relationship.
Consult legal counsel for your specific circumstances.*