Examining AI-Powered Mortgage Through the Lens of Federal Law

Prepared for MBA by Orrick, Herrington & Sutcliffe, LLP

Special thanks to Jeffrey P. Naimon, Caroline Stapleton, Steve vonBerg, and Lauren Bomberger

Table of Contents

Introduction

Mortgage lenders are increasingly developing and deploying artificial intelligence ("AI") tools, including generative and agentic AI capabilities, across a broad range of operational and governance functions. Indeed, some broader definitions[1] of the term "AI" encompass algorithmic decisioning tools like Desktop Underwriter and Loan Prospector that have been used by mortgage companies for decades.

In recent years, lenders have turned to more novel AI capabilities for their potential to significantly increase efficiency and accuracy in marketing, applicant engagement, loan processing and underwriting, customer services, servicing functions generally, and legal and compliance reviews and reporting. While the financial services sector historically had been more cautious in embracing novel AI technologies — particularly in the absence of clear regulatory endorsement — sentiment has materially shifted in recent years, with new AI use cases being embraced by mortgage companies of varying sizes and operational complexity.

AI's assistance with — and, in some cases, performance of — a broader range of mortgage-related tasks raises novel questions about expectations for human involvement with AI models, as well as risk management more broadly. For example, some members seek regulatory guidance confirming mortgage loan originator licensing and registration expectations in circumstances where AI is able to independently take loan applications, correspond with loan applicants, assess credit risk, and render a credit decision. The mortgage industry also has yet to coalesce around a uniform, principles-based risk management framework that lenders can adapt based on their own particular circumstances and AI uses.

Informed by discourse with MBA members, this white paper seeks to address a number of these key issues, and to provide actionable guidance for mortgage companies seeking to integrate AI solutions with appropriate oversight and human support.

1. *See, e.g.,* 15 U.S.C. § 9401(3). Note that there is not a universally-accepted definition of "AI" across industries, and different stakeholders may assign divergent meanings to this term.

Overview of AI Adoption and Regulation in the Mortgage Industry

Our members' engagement with AI has continued since MBA's last white paper exploring AI use cases in the mortgage industry more than a year ago.[2] Mortgage companies are frequently deploying and, increasingly, developing their own AI tools to maximize customer engagement, increase efficiency in the lead development, application, and origination processes, and improve credit risk determinations to more accurately predict loan performance.

As one of our members observed, "we don't know where AI is going to be in 90 days because it's advancing so fast." The result is a diverse and rapidly expanding ecosystem of AI tools that touch nearly every aspect of the mortgage business.

Notably, the term "artificial intelligence" can have varying meanings depending on the stakeholder. While the lack of a universally-accepted definition can, at times, make it challenging to converse about AI adoption and risk management, there is general acceptance of three main categories of AI: (i) generative AI, which generates new content in the form of text, audio, images, or video in response to a prompt; (ii) predictive AI, which learns from historical data to make predictions about future events; and (iii) agentic AI, which is designed to interact independently with its environment.

In conversations with members, MBA learned about various use cases for generative, predictive, and agentic AI in mortgage lending:

- Many customer relationship management systems now incorporate AI-powered bots and conversational agents. These tools handle inbound consumer inquiries, schedule calls with loan officers, and answer routine questions. These bots and AI agents are also used to respond to servicing requests from consumers.

- Some mortgage lenders are engaging with applicants via chatbots throughout the application process, from responding to initial questions presented by prospective applicants on the lender's website, to requesting and receiving application materials and communicating credit decisions. However, members reported varying levels of comfort with respect to permitting large language models ("LLMs") to communicate with applicants and prospective applicants, as opposed to rules-based chatbots that have become a standard offering in the industry.

- Companies are considering, and in some cases already leveraging, telephonic AI-powered call center representatives to engage in outbound and inbound customer calls. In general, members reported that telephonic agents (i) proactively disclose that they are non-human, AI representatives, and (ii) offer customers the option to speak to a human representative if so desired.

- AI continues to be leveraged in models for credit underwriting, pricing, and fraud detection, either on its own or in combination with rules-based, static models.

- At the most advanced level, members reported that agentic AI systems may soon be capable of conducting end-to-end mortgage application and origination activities. This includes receiving a loan application, underwriting the loan, determining eligibility for sale to an investor, and presenting final documents for customer review within hours, rather than weeks. Notably, these tools would be capable of operating in a manner that does not require human participation in, or approval of, a particular transaction prior to completion, regardless of whether these tools are actually used in this way.

2. MBA, AI in the Mortgage Industry: Current Use of AI in the Mortgage Industry and Principles for Lawmakers (Nov. 2024).

This shift toward a more comprehensive integration of AI in mortgage lending is also being recognized by secondary market investors. In March 2026, Freddie Mac updated its Seller/Servicer Guide ("Guide") to introduce additional requirements related to the use of AI and machine learning, including that seller/servicers must establish clear governance frameworks for AI and machine learning adoption.[3] The update also incorporated AI risk management considerations into the Guide's existing information security requirements.[4] While the revisions to the Guide are relatively high-level, they demonstrate Freddie Mac's recognition that companies are, and will continue, using AI in mortgage origination and servicing. Similarly, Fannie Mae published guidance in an April 2026 lender letter calling for seller/servicers to have policies and procedures in place "regarding the development, implementation, and maintenance of any AI/ML system it utilizes, and the measuring and managing of AI/ML risks."[5]

In contrast, federal lawmakers have been less clear as to their expectations for the use and oversight of AI in mortgage lending. There are some guideposts that are useful to mortgage companies, such as (i) the Consumer Financial Protection Bureau's ("CFPB") 2024 rulemaking regarding the use of algorithms and AI in home appraisals and valuations,[6] and (ii) the federal banking agencies' model risk management guidance, amended in April 2026, that sets forth expectations for developing, validating, and monitoring some (but not all) types of AI models. Nevertheless, at the federal level, the industry's use of AI largely continues to be governed by the existing financing and consumer protection laws that have applied to mortgage lenders for decades. While some states, including Colorado, have enacted their own laws governing AI use in consequential decisions impacting consumers, including lending,[7] there is still uncertainty about how these laws will ultimately be interpreted and enforced. In short, legislators and regulators have provided few concrete answers to the questions that mortgage companies have about expectations for, and any limitations on, their use of AI in originating and servicing mortgage loans.

3. Freddie Mac, Single-Family Seller/Servicer Guide, Sec. 1308.8 (rev. Mar. 6, 2026).
4. *Id.,* Sec. 1302.2.
5. Fannie Mae, LL-2026-04, Governance framework on use of artificial intelligence and machine learning (Apr. 8, 2026).
6. Quality Control Standards for Automated Valuation Models, 89 Fed. Reg. 64538 (Aug. 7, 2024).
7. *See, e.g.,* An Act Concerning the Use of Automated Decision-Making Technology in Consequential Decisions, and, in Connection Therewith, Making an Appropriation, Colo. S.B. 26-189 (2026).

Key Legal Questions About the Use of AI in the Mortgage Industry

Mortgage loan originators ("MLOs") play the principal consumer-facing role in the mortgage lending process, including accepting loan applications, negotiating the terms of the mortgage loan, and serving as a primary point of contact for applicants.

Over time, MLOs, along with other mortgage professionals, such as loan processors and underwriters, have increasingly leveraged technology to enhance their speed, accuracy, and compliance with investor requirements. Commonly used tools include automated underwriting systems, text-recognition software, and electronic versions of disclosures, agreements, signatures, and communications. Many of these tools have been utilized in mortgage lending for decades and are essential to the effective functioning of the industry today.

Technological innovations in mortgage origination have been deployed at the direction of human employees. In particular, MLOs generally have had direct personal involvement with each loan application and the corresponding applicant(s). With the introduction of generative and agentic AI, however, it is now possible for tasks previously reserved for human MLOs — such as applicant communications– to be performed exclusively using LLMs and predictive algorithms. Mortgage bankers are increasingly faced with questions about the extent to which models can perform MLO tasks, and what (if any) level of human oversight is required to maintain compliance with applicable laws and regulations.

This section addresses frequently-asked questions from mortgage companies in connection with AI deployment and explains how existing regulatory compliance requirements likely apply to AI tools in the origination and servicing processes.

A. SHOULD AI TOOLS HAVE THEIR OWN MLO LICENSE OR REGISTRATION?

The Secure and Fair Enforcement for Mortgage Licensing Act of 2008 ("SAFE Act")[8] established a nationwide system for MLO licensing and registration. This statute was enacted in response to the mortgage crisis and was designed to enhance consumer protection and reduce fraud by ensuring that MLOs meet minimum standards of competency, character, and fitness, and are subject to regulatory oversight. While it does require human MLOs to be licensed or registered (depending on the nature of their employment), **the SAFE Act does not require AI tools to have their own MLO license or registration to engage in loan origination activities.**

Critically, the SAFE Act defines the term MLO to mean, with certain exceptions, "an individual who (1) takes a residential mortgage loan application and (2) offers or negotiates terms of a residential mortgage loan for compensation or gain."[9] In addition, the SAFE Act prohibits "individual[s]" from engaging in the business of a loan originator without (i) the required license or registration and (ii) a unique identifier.[10] The statute's repeated use of the term "individual" in reference to MLOs and loan origination activities strongly suggests that the SAFE Act's licensing and registration provisions apply to individual human beings, and not to systems, programs, or models. The following provisions provide further support for this interpretation:

8. 12 U.S.C. §§ 5101–5117 (2008).
9. 12 U.S.C. § 5102(4) (emphasis added).
10. Registered or licensed MLOs receive a unique identifier that is permanently assigned to them by the Nationwide Mortgage Licensing System and Registry ("NMLS" or "NMLSR").

- The statute imposes pre-licensure education, examination, background checks (including fingerprint collection), and continuing education requirements on MLOs — which, by their nature, can only be reasonably applied to human beings.

- The definition of a "registered loan originator" is "any individual who… is an employee" of a depository institution or other designated entity. At present, there is no precedent finding that anything other than a human person — including an AI system or model — can be an employee.

- Applicants for MLO licensure must meet "character" requirements, another concept that has not — at least, not yet — been applied to AI systems or models.

Based on the plain language of the SAFE Act and the practical application of its requirements, it is clear that Congress intended the MLO licensing and registration regime to apply to human individuals, and that the statute cannot reasonably be expanded to encompass AI tools used by mortgage companies in the loan origination process. Similar language defining the term "mortgage loan originator" as being an individual — and, thus, presumably a human — is found in state licensing and registration laws.[11] We recognize that, from a practical standpoint, lawmakers could not have reasonably contemplated the extraordinary capabilities that AI would develop over the ensuing decades. Still, the statutory language would require legislative amendment in order to expand licensing and registration requirements to AI systems or models themselves.

It is important to note that mortgage companies that deploy AI systems in the origination process remain subject to applicable entity-level licensing and regulatory requirements. For example, a mortgage company must be licensed as a mortgage lender or mortgage broker in the jurisdictions in which it operates and should ensure compliance with all applicable federal and state lending and consumer protection laws. Moreover, the individual employees of the mortgage company who supervise, manage, or oversee an AI system may themselves, depending on the nature of their activities, be required to be licensed or registered as MLOs. For example, if a human employee reviews and communicates credit decisions generated by an AI system to an applicant, or communicates loan terms generated by AI to a consumer, that employee may be engaged in loan origination activities requiring licensure.

11. *See, e.g.,* Cal. Fin. Code § 22013(a) ("'Mortgage loan originator' means an individual who, for compensation or gain, or in the expectation of compensation or gain, takes a residential mortgage loan application or offers or negotiates terms of a residential mortgage loan.") (emphasis added).

B. CAN MORTGAGE LOANS BE ORIGINATED WITH NO HUMAN INVOLVEMENT?

The SAFE Act makes clear that if a human is engaged in loan origination activities, that person must be licensed or registered as an MLO, with a unique identifier assigned by the NMLS ("NMLS ID"). However, the SAFE Act does *not* answer the broader question of whether mortgage companies can offer completely human-free loan originations. As explained below, for consumer credit transactions under the Truth in Lending Act ("TILA")[12] and Regulation Z,[13] **a human MLO should still be assigned to each mortgage loan origination, and their NMLS ID must be disclosed to the applicant on the credit application and other disclosures and documentation provided to the consumer.**

While there is no express federal mandate that a human MLO must participate in the mortgage origination process, TILA and Regulation Z effectively impose this requirement through disclosure mandates. Specifically:

- Regulation Z defines the term "loan originator" to include "individual loan originators" — i.e., natural persons — and "loan originator organizations" which includes any loan originator that is not an individual loan originator and typically refers to a mortgage broker company.[14]

- For consumer mortgages, lenders must include "[t]he name of the individual loan originator (as the name appears in the NMLSR) with primary responsibility for the origination and, if the NMLSR has provided such person an NMLSR ID, that NMLSR ID" on the credit application, the Loan Estimate and Closing Disclosure, the note or loan agreement, and the security instrument.[15]

- The Loan Estimate is also required to disclose "[t]he name and NMLSR ID of the individual loan officer (labeled "Loan Officer" and "NMLS ID/License ID," respectively) of the creditor and the mortgage broker, if any, who is the primary contact for the consumer."[16] The commentary to Regulation Z states that 12 C.F.R. § 1026.37(k)(2) "requires the disclosure of the name and NMLSR ID of the person who is the primary contact for the consumer, labeled 'Loan Officer.' The loan officer is generally the natural person employed by the creditor or mortgage broker disclosed under [Section] 1026.37(k)(1) who interacts most frequently with the

12. 15 U.S.C. §§ 1601–1667f (1968).

13. 12 C.F.R. Part 1026.

14. *Id.* § 1026.36(a)(1).

15. *Id.* § 1026.36(g) (emphasis added).

16. *Id.* § 1026.37(k)(2) (emphasis added).

consumer and who has an NMLSR ID…"[17] This comment appears to assume that there will be a human MLO, with a name and NMLS ID, who acts as the primary point of contact for a mortgage loan applicant.

Taken together, the disclosure requirements of Regulation Z implicitly mandate the assignment of an individual MLO to each mortgage loan transaction. From a practical standpoint, federal and state licensing regimes appear to have been structured with an underlying assumption that an individual MLO would be involved in each mortgage origination, and that mortgage companies would employ human MLOs to conduct their business. Informal guidance from the Texas Department of Savings and Mortgage Lending, for example, explains the requirements for entity-level registration by stating that "[t]he entity holding the mortgage company license or mortgage banker registration acts by and through individual RMLOs that are licensed by the Department. As a result, an entity license or registration is required to conduct residential mortgage loan origination business."[18]

Consequently, even where an AI system performs most of the loan origination tasks, the consumer should be provided the name and NMLS ID for a human MLO assigned to their application. The AI tool should also clearly disclose at the outset of communications with the consumer, and in any subsequent communications, that it is not a human and/or that it is generative AI.[19]

C. WHAT ROLE MUST A HUMAN MLO HAVE IN THE MORTGAGE ORIGINATION PROCESS?

While Regulation Z effectively requires that a human MLO with an NMLS ID be identified as the MLO with "primary responsibility" or as the "primary contact" for each consumer mortgage transaction, it does not shed light on the level of actual involvement that a human MLO must, or should, have in the mortgage origination process. In the absence of a clear legal standard on this point, courts and regulators instead could look to general consumer protection statutes at the federal and state levels that prohibit misrepresentations in consumer lending transactions. For example, **identifying an MLO on a consumer's mortgage loan application, when in fact the entire origination will be conducted by AI, could be viewed as a deceptive act or practice —** though a consumer-facing disclosure stating that the application process is performed AI only, with no human involvement, could mitigate this risk.

Section 5 of the Federal Trade Commission Act ("FTC Act") prohibits unfair or deceptive acts or practices ("UDAPs"),[20] and Section 1036 of the Dodd-Frank Wall Street Reform and Consumer Protection Act extended this prohibition to include abusive acts or practices for financial services providers within the CFPB's jurisdiction.[21] In addition, virtually all states have enacted analogous consumer protection statutes, often referred to as "mini-UDAP" statutes, which impose similar or even broader prohibitions on deceptive and unfair business practices.

FTC guidance suggests that a lender's representation is "deceptive" if (i) it is likely to mislead a consumer, (ii) the consumer is acting reasonably under the circumstances, and (iii) the representation is material.[22] Importantly, it is not always necessary for a court or regulator to conclude that a consumer was actually harmed in order to find that a deceptive act or practice took place — only that injury was likely.[23] With respect to AI-based loan origination:

17. Comment 3 to 12 C.F.R. § 1026.37(k).

18. Texas Dep't of Savings and Mortgage Lending, FAQs: Licensing/Registration of Mortgage Loan Origination Entities (emphasis added).

19. Utah has already enacted legislation addressing AI communications with consumers in regulated industries and has adopted such disclosures as a safe harbor for compliance. *See* Utah Code Ann. § 13-77-104. Other states could foreseeably interpret their fair trade practices laws to require such disclosures

20. 15 U.S.C. § 45.

21. 12 U.S.C. §§ 5531, 5536.

22. FTC Policy Statement on Deception, 103 F.T.C. 110, 174 (1984).

23. *Id.*

- Disclosing to a consumer that a specific MLO is the "primary contact" for their mortgage application arguably creates a reasonable expectation that the named MLO will, in fact, be available to the consumer, have some involvement in their application and be able to answer questions from the consumer about their application.

- If the named MLO has no actual involvement with the consumer's loan because the origination process is conducted entirely using AI tools, then the company's disclosure could be viewed as "likely to mislead" a reasonable consumer.

- If the MLO does perform some level of review or have some customer interaction, this could reduce — but not entirely eliminate — the UDAP risk. If the consumer reasonably believes that there is a human MLO acting as the "primary contact" for their application, but this individual is not actually available to discuss their loan file, a court or regulator could still conclude that the lender misrepresented the MLO's role in the application.

- A lender performing application intake, underwriting, and loan origination exclusively using AI, with no human involvement, may be able to mitigate this UDAP risk by disclosing this fact to consumers at appropriate points in the origination process.

Beyond federal and state UDAP risk, sound risk management practices also counsel in favor of a meaningful role for human MLOs in the origination process. AI systems, regardless of their sophistication, can produce erroneous or unexpected outputs, including hallucinations, misinterpretations of borrower circumstances, or failure to consider relevant information that a trained, experienced human MLO would recognize. Ensuring that the individual MLO assigned to a particular application is familiar with its contents, has confirmed that any key AI outputs are appropriate, and/or is available to serve as a point of contact for the applicant can serve as a critical check against these risks.

Accordingly, while there currently are no firm standards for MLO involvement in mortgage loan applications and origination, companies should consider the following guideposts as they increasingly leverage AI tools for these activities:

- Establish internal-facing expectations for MLO involvement in each loan application to which their unique identifier is assigned. Expectations may depend, for example, on the complexity of the loans, the degree of AI autonomy, the nature of the MLO's oversight responsibilities, and the experience of the MLO. The guiding principle is that the MLO should have at least some level of involvement in each loan application, even if one or more AI systems play a significant role.

- The human MLO disclosed on the application and other loan-related documents should be available to the borrower as the primary point of contact. Chatbots and other LLM-based communication channels with applicants should provide the ability to "off-ramp" to the MLO at the applicant's request.

- Perform pre-launch and periodic ongoing validation of AI models used to perform origination-related activities, to confirm that they are still accomplishing the purpose(s) for which they were developed, and that applicant outcomes are consistent with, or better than, the outcomes that the company would have expected had the entire process been conducted by humans.

Ultimately, mortgage companies should assess their own risk tolerances when determining the extent of human MLO involvement in the mortgage loan application and origination process. In particular, disclosures that the consumer is dealing with an AI system may not be sufficient to relieve a lender of its obligation to comply with Regulation Z, or to fully mitigate UDAP risk. Even if an applicant is aware that their application is effectively reviewed, processed, and underwritten entirely using AI, a court or regulator could still find that the applicant was entitled to review of their application by a human MLO.

D. HOW CAN MORTGAGE COMPANIES MANAGE RISK WHEN INCORPORATING AI TOOLS INTO SERVICING?

While much of the progress made in AI innovation in the mortgage industry has been focused on application processing and loan origination, mortgage loan servicers are also preparing (and, in some cases, have already begun) to deploy AI-powered technologies for managing incoming servicing-related inquiries, administrative tasks, loss mitigation, and others. Regulation of mortgage servicing is highly prescriptive and rules-based and, thus, the main risks that servicers face relate to non-compliance with the mortgage servicing rules set out in Regulation Z and Regulation X,[24] which implements the Real Estate Settlement Procedures

24. 12 C.F.R. Part 1024.

Act ("RESPA"),[25] as well as the Fannie/Freddie Servicing Guides, FHA Handbook, VA Servicing Guide, and bespoke investor requirements in the secondary markets.[26]

The complexities associated with mortgage serving are particularly pronounced in connection with loss mitigation. Indeed, in a 2023 white paper, MBA highlighted the need for policymakers to focus on promoting simplicity, standardization, and sustainability in loss mitigation to ensure servicers can continue to effectively work with struggling homeowners.[27] Three years later, members are exploring ways in which AI tools may be well-positioned to help servicers navigate this complexity and to facilitate communications with borrowers regarding their options in the event of nonpayment or default. **Lenders should ensure that any AI tools used in servicing are validated pursuant to their risk management policy and procedures, and should consider including a human in the loop for higher-risk decisions or interactions relating to loss mitigation and/or foreclosure.**

To the extent that generative AI technology is used for phone calls or other borrower communications, it is a best practice (and emerging legislative standard[28]) to disclose in each communication session that the consumer is not interacting with a human, but rather generative AI. There are also additional considerations under the Telephone Consumer Protection Act ("TPCA"), which restricts the use of "artificial or prerecorded voice."[29] In 2024, the Federal Communications Commission ("FCC") issued a declaratory ruling that this would encompass current AI technologies that generate human voices and, as such, calls that use AI would require prior express consent.[30]

25. 12 U.S.C. §§ 2601–1617 (1974).

26. Freddie Mac, Single-Family Seller/Servicer Guide; Fannie Mae, Single Family Servicing Guide; HUD Handbook 4000.1, FHA Single Family Housing Policy Handbook; VA Loan Electronic Reporting Interface, Servicer Guide.

27. BA, "The Future of Loss Mitigation" (Feb. 2023).

28. *See, e.g.,* Utah Code Ann. § 13-77-104.

29. 47 U.S.C. § 227(b); 47 C.F.R. § 64.1200(a)(1), (3).

30. FCC 24-17 (Feb. 8, 2024).

Developing a Framework for AI Risk Management in Mortgage

A consistent theme emerging across industries is the need for unified, industry-specific governance frameworks for the use of AI, where the technology is supplanting human activity that has historically been subject to heightened regulation. The concern is straightforward: in the absence of an industry-developed shared framework, federal and individual state regulators will be incentivized to fill the void and may adopt their own requirements that result in a patchwork of costly, confusing, and potentially contradictory compliance obligations. As one MBA member remarked, if the industry fails to act collectively, "fifty different states will regulate us in fifty different ways."

As noted previously, some states have already begun to legislate. Colorado has enacted legislation aimed at prohibiting algorithmic discrimination in consequential decisions using automated decisioning technology, including lending,[31] and other states are likely to follow suit. Still, there remains some hesitation among legislators when it comes to enacting proscriptive requirements as they balance the need for appropriate controls with a desire to foster innovation. The mortgage industry may have an opportunity to get ahead of likely-forthcoming AI legislation by coming together to propose a principled, risk-based framework that regulators and legislators could potentially adopt outright or use as a safe harbor, or other reference point, for demonstrating good-faith compliance with risk management expectations.

In addition to the above, creating a framework can be useful to establish a starting point for industry participants in thinking about how to manage AI risk within their own institutions. The framework should be sufficiently broad to apply to the entire industry while also enabling individual companies to tailor the guidelines to particular use cases and risk tolerances. For example, a formal model risk management program may not be practical or necessary for an independent mortgage bank ("IMB") that is not subject to prudential regulation. Instead, the IMB could rely on the framework to identify core risks and controls (some of which are discussed below) that are relevant to their individual risk profiles and tolerances, and use these to adopt their own policies, procedures, and best practices for AI risk management.

A. EXISTING FRAMEWORKS AS A FOUNDATION

Rather than building a framework from scratch, mortgage industry participants should consider leveraging existing, well-established risk management frameworks. For example, the AI Risk Management Framework developed by the National Institute of Standards and Technology ("NIST") has been identified by regulators and leaders across industries — including financial services[32] — as a strong foundation for identifying, understanding, and controlling for AI-related risks. Within financial services, there are a number of additional reference points, including federal prudential regulatory guidance on model risk management, interagency cybersecurity frameworks such as the FFIEC Cybersecurity Assessment Tool, and a recently-published, NIST-based framework for financial institutions developed by the U.S. Department of the Treasury in partnership with a variety of public and private groups and organizations. Most relevant, the Mortgage Industry Standards Maintenance Organization ("MISMO") recently issued its Framework for Responsible Artificial Intelligence in Mortgage Ecosystems

31. *Supra* n. 7.

32. See the Financial Services AI Risk Management Framework, published by the Financial Sector AI Executive Oversight Group in February 2026.

("FRAME"), which integrates key aspects of the NIST framework and provides MISMO members with useful guidance and instruction regarding the risk assessment and oversight of AI use within their organizations.

A principles-based, risk-based approach is essential to accommodating the wide range of company sizes, business models, and technological capabilities within the mortgage industry, and avoids the problem of prescriptive rules becoming obsolete as technology advances.

B. STRUCTURE AND CONTENT OF THE FRAMEWORK

I. DEFINITIONS AND SCOPING

The framework should provide clear, consensus definitions of key terms, including "artificial intelligence," "predictive AI," "generative AI," and "agentic AI." The lack of consistent definitions is a significant barrier to consistent risk management among mortgage companies. Indeed, industry participants have noted that some regulatory frameworks, including those published by government-sponsored enterprises, do not even define what AI is. Conversely, other frameworks define AI so broadly that the term would include algorithm-based decision-making tools like Desktop Underwriter and Loan Prospector that have been in use for decades, and which were never viewed as "artificial intelligence" warranting heightened risk management. An industry-specific definition may be useful to distinguish between machine learning and other novel AI tools and long-established models, software, and systems.

Some existing frameworks, such as the federal prudential banking regulators' Supervisory Guidance on Model Risk Management ("Interagency MRM Guidance"), define their scope to exclude certain types of AI. Specifically, the Interagency MRM Guidance — amended in April 2026 — states that because generative and agentic AI models are "novel and rapidly evolving," they are not included within the scope of the guidance document. However, the agencies note that these types of AI should still be subject to an entity's risk management and governance practices to determine appropriate governance and controls for generative and agentic AI tools, processes, or systems.[33]

In sum, the proposed framework should make clear what models and systems the term "AI" includes, and whether there are any scoping limitations on the types of AI covered by the framework's recommendations.

II. CLASSIFICATION, INVENTORY, AND RISK ASSESSMENT OF AI USE CASES

The framework should include a classification matrix that categorizes the principal ways AI is used in mortgage, facilitating each entity's assessment of (i) the level of risk presented by such uses, and (ii) the appropriate level of oversight for each. Use cases could include, for example, (i) consumer-facing AI tools that conduct activities that would otherwise be performed by a licensed human individual (highest oversight); (ii) consumer-facing chatbots and voice agents that handle general inquiries (high oversight); (iii) operational efficiency tools such as machine learning-based income calculators and document processors (moderate oversight); and (iv) internal administrative tools used for tasks such as policy review and gap analysis (lower oversight).

The framework should additionally call for lenders to compile and maintain a comprehensive inventory of AI tools or models. Inventories should reflect each tool or model's level of complexity, usage, risks, and interactivity with, or dependencies on, other models.[34]

Once AI tools and models are inventoried and classified, they can be subject to risk assessment and tiering. The framework may leverage existing AI risk management frameworks, such as the framework published by NIST, as a starting point for developing risk assessment procedures, and then customize the framework to the industry's needs by incorporating mortgage-specific risks, legal requirements, regulatory guidance, and operational considerations.

III. GOVERNANCE AT THE ENTERPRISE LEVEL

The framework should address which individuals and/or committees within an organization should bear primary responsibility for AI oversight.[35] Industry participants have suggested that for AI systems conducting licensed activities or making credit decisions, board-level and C-suite approval may be appropriate. For lower-risk AI deployments, such as internal administrative tools, department-level oversight may be sufficient. The framework should additionally set the expectation that appropriate AI governance requires tailored policies and procedures to ensure that responsible individuals understand and are able to effectively carry out their oversight roles.[36]

33. OCC, FRB, and FDIC, Supervisory Guidance on Model Risk Management (Apr. 17, 2026). (Last accessed May 26, 2026.)

34. *See e.g.* Interagency MRM Guidance, at 11.

35. *See e.g.* id., at 10 ("Model risk management benefits from clear roles and responsibilities with well-defined accountability, including with respect to potential conflicts of interest (e.g., misalignment of incentives between different reporting lines, such as model development and validation groups). Sound governance practices delineate the individual(s) responsible for key activities throughout the model lifecycle, from development through validation and ongoing monitoring.").

36. *Id.*

IV. MODEL REVIEW AND VALIDATION

Among the most important decisions mortgage companies must make when using AI is how frequently AI models should be reviewed and validated.[37] Traditional model validation cycles — often annual — may not be sufficient for AI models, particularly those that are capable of self-learning and updating their own parameters. If an AI model is capable of changing its own parameters, or will be subject to a human-initiated update more than annually, it should be reviewed at more frequent intervals. If, by contrast, the model is static, traditional review cycles may be adequate. In all cases, mortgage companies should maintain version control documentation that catalogs changes to AI models, including changes the model has made to itself, and that allows the company to identify and assess the impact of those changes.

The practical tension inherent in AI model review should be acknowledged. If the review and oversight requirements are too onerous, the efficiency gains that AI provides will be negated, undermining the business case for AI adoption. The goal of the risk framework is to give companies the ability to calibrate their approach based on the risks presented by each specific AI use case and/or model, with the goal of ensuring meaningful oversight without eliminating the benefits of AI.

37. *Id.* at 7 ("Validation provides insight into the reliability of a given model, based on its underlying assumptions, methods, data, and relevant theories, as appropriate. This understanding helps characterize the source and extent of the model risk.").

C. KEY RISKS TO CONSIDER

I. FAIR LENDING AND BIAS

The use of AI in mortgage lending implicates a number of federal and state fair lending laws, including the Equal Credit Opportunity Act ("ECOA")[38] and the Fair Housing Act ("FHA").[39] AI systems, including the most sophisticated LLMs, are trained on historical data. If that data reflects historical patterns of discrimination, there is concern that the AI system may perpetuate or amplify those patterns in its outputs. There is also particular concern about the speed at which AI can make decisions. For example, a discriminatory pattern that might be identified and corrected over weeks or months of human decision-making could be replicated thousands of times in a matter of hours by an AI system.

An effective risk framework in mortgage lending requires robust fair lending compliance and testing protocols. These protocols may appropriately include periodic statistical analyses of AI-generated outcomes across protected classes, testing for disparate impact and disparate treatment, and comparison of AI outcomes against human decision-making benchmarks. The frequency and rigor of testing should be calibrated to the volume of decisions being made and the speed at which the AI system is operating.

II. EXPLAINABILITY

Under ECOA and its implementing regulation,[40] Regulation B, a creditor that takes adverse action on a credit application must provide the applicant with a notice of adverse action that includes a statement of the specific reasons for the action taken.[41] In recent years, regulators have expressed concern that the use of AI in credit decisioning has the potential to complicate or undermine lenders' compliance with this requirement. Notably, AI models may consider hundreds or thousands of variables in reaching a credit-related decision or recommendation, and providing only four reasons for adverse action may not accurately reflect the collective reasons for an unfavorable outcome. Moreover, AI models may learn and change over time, such that the same inputs could produce a different outcome if submitted at a different time. Prior CFPB guidance, which has since been rescinded, emphasized the importance of both accuracy and specificity when stating the specific reasons for the action taken.[42]

38. 15 U.S.C. §§ 1691–1691f (1974).

39. 42 U.S.C. §§ 3601–3619 (1968).

40. 12 C.F.R. Part 1002.

41. 12 C.F.R. § 1002.9.

42. CFPB, Circular 2023-03 (Sept. 19, 2023) (rescinded).

A practical framework should address explainability and establish standards for providing effective and compliant adverse action notices when using AI models in credit-related decisions. For example, the industry may agree that for decisions made by models leveraging large quantities of variables, it is appropriate to use Shapley values to determine which variables had the greatest impact on the decision.[43] The framework should also consider the impact of human review overlays, and how notice content could be impacted where a hybrid of AI model outputs and human underwriting decisions collectively result in a denial or other adverse action.

III. STEERING

AI systems that interact directly with consumers and are capable of advising consumers on available loan products and pricing may make recommendations based on criteria that are generally viewed as problematic if the same recommendations were made by a human MLO. Regulation Z prohibits a human MLO from directing or "steering" a consumer to consummate a transaction based on the fact that the MLO will receive greater compensation for that transaction than in other transactions the MLO could have offered to the consumer. The potential exists for an AI system to recommend a mortgage product to a consumer based on the likely gain on sale or other revenue to be realized by the mortgage company rather than based on the applicant's interests. Even though Regulation Z's prohibitions on steering apply only to loan originators (i.e., humans, as discussed above), the CFPB could still determine that such actions by an AI system constitute an unfair, deceptive, or abusive act or practice.

Mortgage companies deploying consumer-facing AI tools should implement guardrails that restrict AI from making recommendations based on loan profitability for the company or otherwise create situations where the applicant is steered towards certain transactions for reasons other than the applicant's stated intent and creditworthiness.

IV. DATA PRIVACY AND SECURITY

The deployment of AI systems, particularly consumer-facing chatbots and voice agents, introduces potential data privacy and security risks. One frequently cited scenario involves a chatbot that has been designed not to solicit personally identifiable information ("PII") from a consumer, but the consumer proactively enters PII into the chat. Companies should be prepared to explain the use, storage and/or disposition of this information, and whether it is used to train models.

Mortgage companies deploying consumer-facing AI tools should implement robust data governance protocols that address, at a minimum, the following: (i) what data the AI system is authorized to collect and store; (ii) how consumer data is used, including whether it is used to train or refine the AI model; (iii) what security controls protect the data; (iv) how long the data is retained; and (v) how consumers are notified of the collection and use of their data. Companies building internal AI tools that access their loan origination systems or business intelligence platforms should also establish clear security frameworks governing the AI system's access to sensitive data.

V. THIRD-PARTY AND VENDOR MANAGEMENT

Many mortgage companies may deploy AI systems developed by third-party vendors rather than building proprietary systems in-house. This raises important questions about the depth of due diligence a mortgage company must conduct on its AI vendors, including the extent to which the company must understand the AI vendor's own vendor relationships (so-called "fourth-party risk"). Members have raised the question of whether a certification process would be appropriate, such as a mechanism by which AI vendors could certify compliance with industry standards, reducing the due diligence burden on individual mortgage companies. Mortgage companies should also consider the implications under the Gramm-Leach-Bliley Act ("GLBA") and Regulation P[44] and their associated disclosure obligations for consumers whose personally identifiable information is input into, or shared with, a third-party AI system. This includes inquiring about and monitoring third parties' data retention and usage policies.

43. The Shapley value is a concept that, in relation to automated models, helps determine the individual marginal contribution of each variable used in the model to the model's predictions. A 2023 white paper published by FinRegLab provides the following explanation: "in a cooperative game with N players and a function that values how much total output is generated if all the players contribute together, the Shapley value is a method that attempts to measure the individual contribution of player i to the output generated by the cooperation of all players." FinRegLab, *Machine Learning Explainability & Fairness: Insights from Consumer Lending* 25 n.13 (July 2023).

44. 15 U.S.C. §§ 6801–09; 12 C.F.R. Part 1016.

Best Practices for Mortgage Lenders

A. DEVELOP INTERNAL EXPECTATIONS FOR A "HUMAN IN THE LOOP"

As discussed above, and consistent with disclosure requirements under Regulation Z, mortgage companies should ensure a human MLO is assigned to each application, and determine the extent to which the MLO will remain "in the loop" for all AI-directed transactions.

B. DEVELOP AN AI RISK ASSESSMENT THAT REVIEWS EACH MODEL OR USE CASE

Mortgage companies should develop classification matrices that identify each AI system in use, the function it performs, the level of risk it presents, and the corresponding level of governance and oversight. Consumer-facing AI systems that conduct activities that would trigger licensing requirements for a natural person, or that make credit or other consequential decisions, should be subject to the highest level of testing, validation, and oversight. In contrast, lower-risk AI deployments may be appropriately managed by the business line, with second line review on an ad-hoc basis. Risk assessment should also consider whether the tool is developed internally or offered by a third party — and, if the latter, should take into account the level of diligence performed on the vendor to confirm the AI tool has been appropriately trained and validated by its developer.

C. IMPLEMENT ROBUST FAIR LENDING TESTING

Mortgage companies using AI in any aspect of the credit decision-making process should ensure they have developed and implemented appropriate fair lending testing protocols, including statistical analysis for disparate impact and disparate treatment, and should maintain documentation sufficient to demonstrate compliance. The speed at which AI operates and self-adjusts may require a testing cadence that is more frequent than what might be adequate for human decision-making or non-machine learning processes. In setting the frequency of testing, companies should also consider their risks commensurate with their size and use of AI.

D. ESTABLISH MODEL REVIEW PROTOCOLS CALIBRATED TO THE AI'S CAPABILITIES

AI models that are capable of self-learning and changing their own parameters should be reviewed at more frequent intervals than traditional models. All AI deployments should be supported by version control documentation that catalogs changes to the model, including changes the model has made to itself, and which allows the company to identify and assess the impact of those changes.

E. NEWER ENTRANTS SHOULD CONSIDER AI-POWERED FRAUD PREVENTION

Fraud prevention represents an immediate, high-value, and comparatively less controversial use case for AI in financial services. AI tools are well suited to detecting patterns indicative of fraud that human reviewers might miss, and the speed and scale at which AI can process data make it a powerful tool for identifying suspicious transactions and detecting anomalies in real time. By reducing fraud losses and associated litigation costs, AI-powered fraud prevention tools can ultimately reduce costs for consumers.

F. PREPARE FOR EVOLVING REGULATORY EXPECTATIONS

As regulatory legislation, guidance, and enforcement priorities continue to develop, mortgage companies should proactively establish internal policies and procedures governing the use of AI that are sufficiently flexible to accommodate new requirements as they emerge. Companies should ensure that their compliance management systems are equipped to address AI-related risks, including through updates to quality control programs, internal audit processes, and consumer complaint monitoring.

G. SEEK OPPORTUNITIES TO ENGAGE WITH STATE LEGISLATORS AND REGULATORS

Once created, the mortgage industry-specific risk management framework discussed above could be previewed for state lawmakers to show that mortgage lenders are proactively and appropriately managing their use of AI. Proactively educating lawmakers and policymakers can benefit mortgage lenders by encouraging practical legislation and regulation and promoting the adoption of requirements that are consistent with practices developed collaboratively by industry participants. Industry members can also enhance the success of advocacy work by highlighting the benefits of AI for consumers, such as faster decision-making and response times.

A key concern for state legislatures and regulators appears to be the expectation that consumers will have access to human review of their loan application decision. Mortgage companies can leverage the framework to advocate for individual MLOs as humans-in-the-loop, rendering additional human review requirements unnecessary in the context of mortgage lending. Legislators' concerns about recourse or liability can be addressed by licensing laws and UDAAP enforcement, much in the same way as pre-AI mortgage loan transactions.

H. ADVOCATE FOR LEGISLATION AND REGULATORY GUIDANCE

The industry should consider whether it would be beneficial to have updated legislation from Congress and state regulators and/or guidance from the CFPB, prudential regulators, and other state and federal agencies on the application of existing consumer protection laws to AI-enabled mortgage processes, including guidance on adverse action notices, fair lending, and the level of human involvement expected in AI-enabled originations. At the federal level, the prevailing view appears to be that existing law is sufficient to regulate AI in mortgage lending; whereas at the state level, we are seeing an increasing number of states attempt to legislate AI use in the financial services industry in different ways. Industry participants should identify the areas where clarity would be useful to allow companies to confidently proceed with deploying AI technology and should also consider whether to advocate for a single unified legislation/guidance, separate federal and state versions, or something else entirely.

Conclusion

The rapid adoption of AI across the mortgage industry presents both significant opportunities and complex legal and regulatory challenges that demand thoughtful, proactive engagement from industry participants.

While existing laws such as the SAFE Act, state licensing requirements, TILA, RESPA, and federal and state UDAP statutes were not written with AI in mind, they nonetheless impose meaningful requirements on mortgage companies deploying AI tools in origination and servicing, particularly with respect to MLO involvement, consumer disclosures, fair lending, data privacy, and others. The absence of comprehensive federal and state guidance on AI in mortgage lending creates an imperative for the industry to develop and adopt a unified, principles-based risk management framework that provides regulators with confidence in the industry's self-governance while preserving the efficiency and innovation that AI enables.

Acknowledgments

MBA would like to acknowledge and thank all the members who have contributed to this workstream and helped offer incredible insight into AI's adoption into the mortgage process.

Special thanks are due to the Chair and Vice-Chair of MBA's Residential Board of Governors during this effort: Dan Sugg, CMB, and Shayna Arrington.

For more information on MBA's AI initiatives, contact:

Gabriel Acosta
Senior Regulatory Specialist,
Assistant Regulatory Counsel
Mortgage Bankers Association
gacosta@mba.org

Liz Facemire, CMB
Associate Vice President, State Government Affairs
Mortgage Bankers Association
lfacemire@mba.org

Richard Hill
Vice President, Information Technology
Mortgage Bankers Association
rhill@mba.org

MBa®

MORTGAGE BANKERS ASSOCIATION

1919 M STREET NW, 5th FLOOR
WASHINGTON, DC 20036