

Insights

DOES YOUR CURRENT USE OF AI IN FINANCIAL SERVICES ALIGN WITH THE U.S. “AI BILL OF RIGHTS”?

Feb 09, 2023

SUMMARY

As OpenAI’s release of ChatGPT in late 2022 and expected release of GPT-4 in 2023 continues to garner widespread attention, there is renewed focus on both opportunities and risks presented by the use of artificial intelligence (“AI”). With this focus comes the inevitable call for regulation. At the end of 2022, the U.S. White House weighed in through what it calls an “AI Bill of Rights” for the American public, a non-binding policy document. Banks and others in financial services should take note of the particular civil rights, privacy, and other priorities expressed in this vision for the future of AI governance.

In financial services, technologies deploying some element of AI are expected to increase but already abound. Just a few examples of such technologies include certain credit underwriting, fraud prevention, anti-money laundering, appraisal, customer experience (chatbots and targeted advertising), and employment (including so-called “bossware”) applications.

There are many open questions regarding the degree to and manner in which existing laws and regulations apply to the use of AI, from anti-discrimination to antitrust to intellectual property laws. It is in large part the rapidly evolving scope and nature of AI that is driving these questions—both in financial services and elsewhere, the term defies a simple, uniform definition. For example, the Consumer Financial Protection Bureau (“CFPB”) has initiated rulemaking to implement Dodd-Frank Act standards governing “automated valuation models” or “AVMs.”^[1] Congress defined AVMs in the legislation mandating these rules as “*any computerized model used by mortgage originators and secondary market issuers to determine the collateral worth of a mortgage secured by a consumer’s principal dwelling.*”^[2] Similarly, under a law pending implementation by New York City governing certain uses of AI in hiring, the threshold term “automated employment decision tool” means “any computational process, derived from machine learning, statistical modeling, data analytics, or artificial intelligence, that issues simplified output, including a score, classification, or recommendation, that is used to substantially assist or replace discretionary decision making for

making employment decisions that impact natural persons. The term ‘automated employment decision tool’ does not include a tool that does not automate, support, substantially assist or replace discretionary decision-making processes and that does not materially impact natural persons, including, but not limited to, a junk email filter, firewall, antivirus software, calculator, spreadsheet, database, data set, or other compilation of data.”^[3]

Existing U.S. bank regulatory guidance reflects both broad and targeted concerns when it comes to AI. The Office of the Comptroller of the Currency (the “OCC”), which supervises national banks and trust companies, maintains a self-described risk-based supervision model in this regard, citing explainability, data management, privacy and data security, and third-party risk as key issues in this space.^[4] In May 2022, CFPB published a circular emphasizing that the adverse action notice requirements of the Equal Credit Opportunity Act (“ECOA”) and its implementing rules, which prohibit discrimination against credit applicants, “apply equally to all credit decisions, regardless of the technology used to make them” and consequently “do not permit creditors to use complex algorithms when doing so means they cannot provide the specific and accurate reasons for adverse actions [such as declined applications].”^[5] In its most recent annual report on credit reporting agency complaints, the CFPB also specifically singled out consumer harm originating from such agencies’ reliance on automated processes to screen and respond to consumer inquiries.^[6]

Through its Office of Science and Technology Policy the White House released broad, aspirational guidance in October 2022 in the form of its “Blueprint for an AI Bill of Rights.”^[7] It also included an accompanying “technical companion” elaborating on the rights it identifies and in general describes these materials together (the “Blueprint”) as embodying a “national values statement and toolkit that is sector-agnostic to inform building these protections into policy, practice, or the technological design process.” This non-binding statement, which purports to reflect feedback from numerous public meetings and other stakeholder engagements, is fairly high-level in nature but does provide significant clues into how AI might be further regulated in the U.S. Institutions and providers in financial services would do well to consider how the priorities outlined in this statement can be accommodated by current AI uses and future plans.

The Blueprint describes the following five (5) fundamental rights:

- **Safe and Effective Systems** - The emphasis here is on pre-deployment testing and other common-sense safeguards against performance failure risks including unintended outcomes and uses. Under this principle, the American public—the intended beneficiary of the White House framework—is “protected from inappropriate or irrelevant data use in the design, development, and deployment of automated systems, and from the compounded harm of its reuse.” As examples of problems associated with inadequate care for this right, the guidance cites an underperforming healthcare model resulting in “alert fatigue” by unreliably flagging a high likelihood of the life-threatening condition known as sepsis, as well as instances of social media platforms silencing users when failing to distinguish between hate speech and such

users' "counter speech" critical of such content. According to the White House, keys to the necessary protection from such outcomes include risk assessment and ongoing monitoring as well as independent testing. Mandatory ethics review and robust data governance are also identified as risk mitigation strategies.

- **Algorithmic Discrimination Protections** - Not surprisingly the White House makes fairness and equity a pillar of its AI "Bill of Rights." Citing risks relating to unjustified disparate treatment and potential violations of existing law, the guidance stresses the importance of proactive and ongoing measures to protect individuals and communities from algorithmic discrimination. This includes an emphasis on representative and robust data—data that is reviewed for bias arising from its historical and social context. This is a reference to the feedback effect seen in models that reflect previous human bias, such as the example included by the guidance in which a hiring tool gave negative weight to resume items using the word "women's" such as "women's chess club captain." Echoing current fair lending law and practice, the guidance recommends that planning and design processes seek to identify systems that result in the least adverse impact to protected classes and communities. It also cites accessibility guidelines and NIST Special Publication 1270, "Towards a Standard for Identifying and Managing Bias in Artificial Intelligence," as resources that can inform how this principle can be put into practice in the context of AI.
- **Data Privacy** - Perhaps the most intuitive element of the White House's framework is data privacy, as it is the combination of vast available input data—often of a sensitive personal nature—and modern computing power that is giving rise to much of AI's accelerating capability. This principle reads like a primer on the EEA's GDPR and other emergent data privacy regimes, calling on increasingly familiar themes such as privacy through design, opt-out rights, rights to deletion, and plain language disclosure as well as freedom from "unchecked surveillance." Here the White House explicitly references the lack of a comprehensive U.S. federal privacy law or regulatory framework and decries the existing "patchwork" of laws in this area and the practices of unscrupulous "data brokers." The Blueprint also bluntly pronounces that traditional terms of use or service—"the block of text that the public is accustomed to clicking through when using a website or digital app—are not an adequate mechanism for protecting privacy." Whether expressed as further foundation for stand-alone federal privacy legislation in the U.S. or for purposes of informing AI-specific laws or private sector design standards, on this front the White House appears to believe that both the gaps and the substantive road map for addressing them are clear. It gives special attention to data relating to sensitive domains, including employment and personal finance.
- **Notice and Explanation** - As a corollary to data privacy considerations, the White House framework conceives of informed consent to the application of automated systems as a fundamental right. The degree of advance notice that is envisioned includes lofty concepts of "interpretability" and "explainability" that should enable those impacted by such systems to

understand and challenge them. Summing this up in due process-like terms, the Blueprint explains: “In order to guard against potential harms, the American public needs to know if an automated system is being used. Clear, brief, and understandable notice is a prerequisite for achieving the other protections in this framework.” The subject of the CFPB’s ECOA adverse action bulletin on complex algorithms, described above, is cited as an existing real-life example of the type of safeguard envisioned here, as it requires lenders to provide reasons for their underwriting and pricing decisions in a manner that allows for error resolution and other remedies. According to the White House, failure in this area leads to automated decision-making systems that are “opaque, complex, and, therefore, unaccountable, whether by decision or omission.” The Blueprint calls for simple explanations of algorithms and other AI technologies that are tailored to the purpose of the particular application at hand, the target audience for the explanation, and the level of risk associated with such application.

- **Human Alternatives, Consideration, and Fallback** - In explaining its final principle the Blueprint asserts that the American public “should be able to opt out from automated systems in favor of a human alternative, where appropriate.” For this purpose appropriateness “should be determined based on reasonable expectations in a given context and with a focus on ensuring broad accessibility and protecting the public from especially harmful impacts.” Such human alternatives would presumably be disclosed in connection with notice and explanation of the system itself, as described above, and would be available to those impacted through a simple, accessible, and timely opt-out. Moreover, the alternative should not be unduly burdensome, and its availability and the human resources devoted to support it should be proportionate to the potential of the automated system to “meaningfully impact rights, opportunities, or access” otherwise afforded to the relevant users or subjects. Interestingly, the White House framework acknowledges that the introduction of manual intervention or support also introduces a special need for attention to training and oversight to mitigate the impacts of human bias. In the customer service context, escalation to a human support team as part of chatbot service or AI-driven call response is cited as an example of this kind of meaningful human alternative.

The bridge in the Blueprint between the White House’s AI “Bill of Rights” and what it describes as a technical companion for its implementation, all as summarized above, is a short section further defining the intended scope of the protections envisioned by the framework. This section defines this scope as encompassing those automated systems that have the potential to meaningfully impact individuals’ or communities’ exercise of (1) civil rights, civil liberties, and privacy, (2) equal opportunities (including to education, housing, credit, and employment), and (3) access to critical resources or services such as healthcare and financial services. This scoping or impact statement presents some parallels to the explicit risk-rating that is at the heart of the EU Artificial Intelligence Act’s pending implementation rules and would apply certain restrictions and documentation requirements on a sliding scale based on risk category. Such an approach would likely be more meaningful and predictable than attempts to regulate AI on a simple defined-term basis alone.

Efforts in the U.S. to develop AI regulation have continued since the release of the White House’s Blueprint for it, with Congress and industry voicing concerns that its approach differs in important ways from the recently finalized NIST AI Risk Management Framework,^[8] and U.S. Representative Ted Lieu even introducing a non-binding resolution generated by ChatGPT calling for Congress to take action to regulate AI.^[9] This discussion is coming at a time when there is heightened attention to the regulation of virtual currency and other blockchain-based applications. These applications fall under elements of current and emerging AI regulation and may—along with certain veins of negative “Big Tech” sentiment—have an outsized effect on additional near-term legislative and regulatory response to this type of innovation generally. Regardless of whether and how such further governmental action unfolds, it is important for institutions, fintech firms, and others within the financial services industry to take stock of current AI uses, prospects, and the early steps of what is likely to be a complex regulatory journey.

[1] <https://www.consumerfinance.gov/about-us/newsroom/cfpb-outlines-options-to-prevent-algorithmic-bias-in-home-valuations/>

[2] https://files.consumerfinance.gov/f/documents/cfpb_avm_outline-of-proposals_2022-02.pdf
(emphasis added)

[3] <https://codelibrary.amlegal.com/codes/newyorkcity/latest/NYCAadmin/0-0-0-135839>

[4] <https://www.occ.gov/news-issuances/congressional-testimony/2022/ct-occ-2022-52-written.pdf>

[5] <https://www.consumerfinance.gov/compliance/circulars/circular-2022-03-adverse-action-notification-requirements-in-connection-with-credit-decisions-based-on-complex-algorithms/>

[6] https://files.consumerfinance.gov/f/documents/cfpb_fcra-611-e_report_2023-01.pdf

[7] <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>

[8] <https://www.nist.gov/itl/ai-risk-management-framework>

[9] <https://lieu.house.gov/media-center/press-releases/rep-lieu-introduces-first-federal-legislation-ever-written-artificial>

RELATED PRACTICE AREAS

- Fintech
- Payment Systems

This material is not comprehensive, is for informational purposes only, and is not legal advice. Your use or receipt of this material does not create an attorney-client relationship between us. If you require legal advice, you should consult an attorney regarding your particular circumstances. The choice of a lawyer is an important decision and should not be based solely upon advertisements. This material may be “Attorney Advertising” under the ethics and professional rules of certain jurisdictions. For advertising purposes, St. Louis, Missouri, is designated BCLP’s principal office and Kathrine Dixon (kathrine.dixon@bclplaw.com) as the responsible attorney.