BY JOHN MULLIN

Artificial Intelligence and Bank Supervision

Regulators are gathering information about how banks use AI

Artificial intelligence has come a long way since English mathematician, logician, and cryptographer Alan Turing's seminal 1950 essay, "Computing Machinery and Intelligence," which explored the idea of building computers capable of imitating human thought. In 1997, almost 50 years after Turing's essay, AI posted a historic breakthrough when the IBM supercomputer Deep Blue won a chess match against reigning world champion Garry Kasparov. Since then, AI's capabilities have improved rapidly, largely through advances in machine learning (ML), especially in ML models that use digital neural networks to classify text, images, or other data. (See "Machine Learning," *Econ Focus*, Third Quarter 2018.) ML is now commonly used in industrial applications, and it underpins a vast number of consumer services, from Google searches to Netflix movie recommendations. Of more recent note, ML technology is the basis of the new generative AI programs, such as ChatGPT, designed to, among other things, conduct useful conversations with human beings.

Financial institutions in the U.S. have hardly sat idle amid these developments. On the contrary, they have developed and implemented AI-based applications for a wide variety of purposes. Yet, overall, the financial industry appears to have taken a gradual approach to AI implementation. McKinsey and Co., in a 2019 survey of the financial services sector, found that only 36 percent of industry respondents reported that their companies had adopted AI for the automation of back-office processes, only 32 percent

had deployed AI-based chatbots for customer service, and only 25 percent had deployed AI for detecting fraud or evaluating creditworthiness. The consulting firm Cornerstone Advisors reported even lower numbers based on its 2022 survey of bank and credit union executives. The firm found that only 25 percent of survey respondents had deployed AI for process automation and only 18 percent had deployed AI-based chatbots.

Whatever the current state of AI deployment in the banking industry, there seems to be little doubt that AI's role in banking has been growing and will continue to grow in importance. Anticipating this growth, U.S. bank regulators continue to monitor and assess banks' use of AI-based applications. In March 2021, the Office of the Comptroller of the Currency (OCC), the Fed Board of Governors, the Federal Deposit Insurance Corporation, the Consumer Financial Protection Bureau, and the National Credit Union Administration issued a request for information (RFI) to improve their understanding of current and prospective bank practices surrounding the new technology. Their efforts are continuing as the technology grows and evolves.

CUTTING COSTS, COUNTERING FRAUD

The most recent generation of chatbots can simulate human conversations and provide bank customers with information on account balances, credit card usage, and interest rates. Capital One, for instance, offers a virtual assistant called "Eno" that can answer client

questions, pay routine bills, and deliver fraud alerts.

Such chatbots may offer benefits to both banks and their clients. For banks, the primary allure may be cost savings. According to a report by consulting firm Deloitte, the top 2,000 U.S. corporations spend roughly $250 billion annually on customer support (50 billion incidents at an average of $5 apiece). For bank customers, much of the upside may come from more rapid and convenient access to information, particularly when that information concerns potentially fraudulent charges against customer accounts.

Nevertheless, many banks appear to be wary of moving too quickly into the realm of automated customer service. Indeed, it appears that the deployment of chatbots has been less common in the banking industry than in other industries. This reluctance may reflect a disconnect between the technology's promise and its present reality. Despite improvements in recent years, surveys show that most consumers still view automated chatbots as sources of great frustration. Banks surely want to cut costs but are naturally hesitant to risk losing long-term customers.

While there are people who may not be looking forward to having more frequent encounters with chatbots instead of live people, some AI applications have been more unambiguously positive for banks and their customers. AI applications using pattern recognition, for instance, have allowed customers to deposit checks online and avoid extra trips to brick-and-mortar bank locations.

AI earns additional high marks for its contribution to fraud prevention.

"Fraud detection is one of the most common uses of AI models in banks, where they have been used for quite a while," says Tom Bilston, an assistant vice president of the Richmond Fed's bank supervisory team and former co-lead of the Fed's Working Group on Artificial Intelligence & Machine Learning. "Credit card fraud is the most common thing that comes up. It can happen when someone acquires a card number and uses it without authorization. But it also happens when people apply for cards using fake identities — this is one place where banks can use AI."

Bank anti-fraud efforts are an escalating game of cat and mouse. "Banks have an interesting reliance on some popular vendor AI solutions and consortium data, given that fraudsters tend to constantly innovate their attack paths," says Susanna Wang, a senior examiner of large financial institutions at the Richmond Fed.

AI technology has also been used by banks to help them comply with their obligations under the 2001 Patriot Act to deter money laundering and the funding of terrorist organizations. "The banking industry's use of AI to uncover unusual payment patterns goes well beyond fraud prevention," says Bilston. "Firms think that AI can help them with their anti-money laundering [AML] and know-your-client [KYC] programs."

Companies such as New York-based Socure have designed identity verification systems that use machine learning to analyze applicants' online, offline, and social data to determine whether they meet KYC standards. Symphony AyasdiAI of Palo Alto, Calif., has developed an AML alert system that uses machine learning to spot suspicious transactions while minimizing the number of false warnings. Data science company Feedzai uses machine learning to help banks monitor transactions; its tool raises red flags when it spots suspicious payment patterns.

AI AND CREDIT EVALUATION

In a matter of more immediate concern for bank supervisors, financial firms have been developing and implementing AI models to support their credit evaluation and loan underwriting processes. "Most of these applications are being developed in the retail space — in credit card and automobile underwriting," says Wang. "For such retail

> "As a general matter, U.S. bank supervisors have found it helpful to think about AI and traditional modeling approaches as being different points on a spectrum rather than as binary possibilities."

applications, the banks must justify their reasoning about credit decisions based on the Equal Credit Opportunity Act, so this is a hotly debated topic of how firms are able to explain their credit underwriting model results when they use AI applications, which are often opaque 'black box' models."

Most, if not all, banks still use traditional credit evaluation models — akin to the models used by the national credit bureaus to calculate consumers' credit scores. (See "Credit Scoring and the Revolution in Debt," *Econ Focus*, Fourth Quarter 2013.) With these traditional models, there is often a single formula used to calculate a credit score based on a relatively small group of indicators, such as an applicant's existing debt service burden and credit history. By contrast, AI models often have multiple layers of complicated analysis involving numerous quantitative and qualitative inputs. As a result, AI models can be much more difficult to understand and interpret than their traditional counterparts.

Bank regulators can leverage existing supervisory guidelines and principles when reviewing banks' use of AI models. In 2011, the Federal Reserve

Board and the OCC jointly issued a document, "Supervisory Guidance on Model Risk Management," to provide banks with comprehensive guidance on how to manage the risks associated with their models, including the potential for adverse consequences due to poor model design or incorrect input data.

"That is generally the framework that banks and regulators point to when thinking about AI," says Wang. "As a general matter, U.S. bank supervisors have found it helpful to think about AI and traditional modeling approaches as being different points on a spectrum rather than as binary possibilities." This approach allows supervisors to bypass the semantic problem associated with defining what is or is not an AI model and to shift the focus toward banks' processes for managing the risks presented by credit evaluation models, whether they are AI or traditional.

The interagency guidance spelled out principles for model design, the monitoring of model usage, and the evaluation of model outcomes. Nevertheless, the guidance recognized that "details of model risk management practices may vary from bank to bank" and placed the ultimate burden on banks to make to maintain "strong governance and controls to help manage model risks."

As a practical matter, bank supervisors do not set out to dictate the particular risk model that a bank should be using. "When we go into a bank with our supervisory lenses, we don't necessarily say something like 'Oh, that algorithm is wrong. You can't use that,'" says Ray Brastow, an economist in the Richmond Fed's Supervision, Regulation, and Credit department. "Our processes are more focused on making sure that the risks associated with a bank's model are being appropriately monitored and controlled."

EXPLAINABILITY

Large AI models based on ML algorithms and trained on large datasets can be largely opaque to humans. While conventional statistical models have well-defined variables and coefficients that experts can interpret, many AI models do not: Under the hood, they're often just a sea of numbers that make up the neural network. Thus, it can be challenging to determine how an AI system arrived at its results. This problem is pervasive across applications that use digital neural networks, including image recognition programs, chatbots, and programs used by scientists to find predictive patterns in fields such as medical research.

Bank supervisors and market commentators are particularly concerned about the potential for AI-based credit models to unintentionally perpetuate human biases such as racism, running afoul of federal anti-discrimination law. In recent years, the Federal Trade Commission and the Consumer Financial Protection Bureau have issued warnings about the potential adverse effects of such "algorithmic biases."

To bank examiners at the Fed and the OCC, the potential for such hidden biases highlights the need for banks to expend the effort and resources necessary to understand the inner workings of their models and to be able to adequately explain model results. "As supervisors, we will evaluate the risks associated with AI models, such as explanatory power, and determine whether the controls are in place to support compliance with applicable laws, rules, and regulations," says the Richmond Fed's Wang.

The OCC's Kevin Greenfield expressed a similar view during his May 2022 testimony before the House Committee on Financial Services, arguing that a lack of model explainability can make it difficult for banks to comply with various regulations, including consumer protection requirements.

The question of explainability was at the top of the list of topics that bank examiners raised in their 2021 RFI, which cautioned that AI systems generally reflect the limitations of their datasets and may "perpetuate or even amplify bias or inaccuracies inherent in the training data."

The consumer advocacy nonprofit Consumer Reports, in its response to the RFI, emphasized the need to safeguard against algorithmic discrimination, arguing, "Claims of objectivity and proof notwithstanding, algorithms can and sometimes do exacerbate bias or have unexpected discriminatory effects, as numerous examples have demonstrated." The organization recommended that credit applicants should be made aware when credit decisions are based on AI algorithms and that such algorithms should be designed with fairness in mind.

In its response to the interagency request for information, the Bank Policy Institute (BPI), which conducts research and advocates for the banking industry, cautioned against excessive requirements for explainability that could stifle innovation. They argued against a one-size-fits-all approach, stressing that explainability should mean different things in different contexts. In their view, it is important to distinguish between explainability in the context of a bank's ability to understand its own models (and describe their workings to supervisors) versus explainability in the context of explaining credit decisions to individual credit applicants. "Consumers want easy-to-understand information on credit decisions," says Chris Feeney, president of BITS, the BPI's technology policy division. "Regulators want explanations and evidence concerning the model architecture and rationale, the sources of data used, the human role in the decision, and the resilience of those models."

Fed Board of Governors then-member Lael Brainard expressed a similar view in a 2021 address, noting, "An explanation that requires the knowledge of a Ph.D. in math or computer science may be suitable for model developers" but a less technical standard may be appropriate in the context of explaining credit decisions to consumers under U.S. consumer protection laws.

The BPI is also concerned that bank regulators may be holding AI-based models to an artificially high standard. "I think one of the concerns is that bank regulators apply stricter standards of explainability to AI models than to standard models," says Paige Paridon, senior vice president and senior associate general counsel of BPI. "There's a concern that banks maybe won't be given the flexibility to experiment with and implement some of these tools and that there's a heightened skepticism coming from bank regulators."

ALTERNATIVE DATA AND MODEL MAINTENANCE

Since AI methods such as machine learning are designed to find patterns by digesting enormous quantities of data, it is hardly surprising that banks would seek out new sources of data to feed into the new models. This possibility, however, has raised concerns in some quarters about the implications of banks' use of "nontraditional" data, which bank supervisors define as information not typically found in consumers' credit files at banks or nationwide consumer reporting agencies. Examples of nontraditional data include information about credit applicants' rent and utility payments as well as the cash flow patterns in their bank accounts.

Bank supervisors issued an "Interagency Statement on the Use of Alternative Data in Credit Underwriting" in 2019 in an attempt to better understand the relevant issues. While recognizing that the use of alternative data has the potential to lower costs and increase credit access, the agencies also pointed out that the use of such data raises questions about how it will affect banks' compliance with consumer protection laws. The 2021

RFI followed up by asking interested parties to provide additional information about their use of alternative data.

Consumer Reports expressed concern with financial firms' control policies with respect to alternative data, particularly in cases where banks may be able to glean sensitive information based on applicants' social media and internet browsing activity: "Not only does this raise privacy concerns that could lead to a chilling effect on free expression, but there is little evidence that these types of data are actually effective in calculating credit risk."

The BPI, in its response to the interagency request, emphasized that the risks of poor data are not unique to AI-based models. Moreover, it pointed out that the data monitoring processes banks use for AI models are consistent with those that they use for their traditional models.

Mortgage lender Quicken Loans, in its response to the RFI, noted that the questions about alternative data are something of a moot point for them, positing that there is little incentive for the mortgage industry to use alternative data sources since they are disallowed by the Federal Housing Agency, the Department of Housing and Urban Development, and the government-sponsored enterprises Fannie Mae and Freddie Mac.

In the interagency RFI, bank supervisors noted their concerns about banks' ongoing maintenance of AI-based credit models, arguing that, since the models evolve over time by "learning" from new data, they may present challenges for model validation, monitoring, and documentation.

Consumer Reports, in its response to the RFI, echoed the concerns of bank supervisors regarding model maintenance, arguing that banks' AI models should be monitored with vigilance to ensure that they do not evolve to incorporate indicators that serve as proxies for prohibited factors such as race.

Supervisors at the Richmond Fed are cautiously optimistic about banks' ability to leverage AI-based models for certain aspects of credit evaluation. "In the banks that we look at, model risk is something they take very seriously," says the Richmond Fed's Brastow. "So even 10 years ago, banks would update their traditional models when they got a bunch of new data. But they didn't just say willy-nilly, 'OK, we're scrapping the old approach.' Instead, banks would evaluate a new model by running it in parallel with its predecessor. And only then, after a lot of time and consideration, would they start making decisions based on the new model, while continuing to run the old model to see how differently the two models perform."

GUARDED OPTIMISM

The U.S. financial sector is still in the early stages of integrating AI into its operations, so there is much anticipation and conjecture as to what will come next. Bank supervisors, while noting many of the potential pitfalls of banks' use of AI-based applications, have conveyed optimism about the technology's potential benefits.

In his 2022 statement before Congress, the OCC's Greenfield emphasized AI's potential to help banks with their regulatory compliance programs, arguing that "AI has the potential to strengthen safety and soundness, enhance consumer protections, improve the effectiveness of compliance functions, and increase fairness in access to the financial services when implemented in an effective manner." He also expressed guarded optimism about banks' use of alternative data, advancing the idea that "alternative data in AI applications may improve the speed and accuracy of credit decisions and may help firms evaluate the creditworthiness of consumers who may not otherwise obtain credit in the mainstream credit system."

Brainard, while concerned about the potential for AI-based credit models to perpetuate biases, pointed to encouraging signs that AI researchers are making some progress toward increasing the transparency of their models, making their results more amenable to explanation. Nevertheless, Brainard argued for caution, stating that "Having an accurate explanation for how a machine learning model works does not by itself guarantee that the model is reliable or fosters financial inclusion. ... The boom-bust cycle that has defined finance for centuries should make us cautious in relying fully for highly consequential decisions on any models that have not been tested over time." **EF**

READINGS

Barefoot, Jo Ann. "The Case for Placing AI at the Heart of Digitally Robust Financial Regulation." Brookings Institution Center on Regulation and Markets, May 24, 2022.

"Global AI Survey: AI Proves its Worth, but Few Scale Impact." McKinsey and Co., Nov. 22, 2019.

"Request for Information and Comment on Financial Institutions' Use of Artificial Intelligence, Including Machine Learning." Office of the Comptroller of the Currency, Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, Bureau of Consumer Financial Protection, National Credit Union Administration, Federal Register Doc. 2021–06577, March 31, 2021.

Shevlin, Ron. "What's Going On in Banking, 2023." Cornerstone Advisors, 2023.