

For release on delivery
10:00 a.m. EST
November 13, 2018

What Are We Learning about Artificial Intelligence in Financial Services?

Remarks by
Lael Brainard
Member

Board of Governors of the Federal Reserve System

at

Fintech and the New Financial Landscape

Hosted by the Federal Reserve Bank of Philadelphia, the Federal Deposit Insurance Corporation, University of Pennsylvania Wharton School of Business, Bank Policy Institute, and Brookings Institution

Philadelphia, Pennsylvania

November 13, 2018

Although it is still early days, it is already evident that the application of artificial intelligence (AI) in financial services is potentially quite important and merits our attention. Through our Fintech working group, we are working across the Federal Reserve System to take a deliberate approach to understanding the potential implications of AI for financial services, particularly as they relate to our responsibilities. In light of the potential importance of AI, we are seeking to learn from industry, banks, consumer advocates, researchers, and others, including through today's conference. I am pleased to take part in this timely discussion of how technology is changing the financial landscape.¹

The Growing Use of Artificial Intelligence in Financial Services

My focus today is the branch of artificial intelligence known as machine learning, which is the basis of many recent advances and commercial applications.² Modern machine learning applies and refines, or “trains,” a series of algorithms on a large data set by optimizing iteratively as it learns in order to identify patterns and make predictions for new data.³ Machine learning essentially imposes much less structure on how data is interpreted compared to conventional approaches in which programmers impose ex ante rule sets to make decisions.

¹ I am grateful to Kelvin Chen and Carol Evans for their assistance in preparing this text. These remarks represent my own views, which do not necessarily represent those of the Federal Reserve Board or the Federal Open Market Committee.

² Executive Office of the President, National Science and Technology Council Committee on Technology, *Preparing for the Future of Artificial Intelligence* (Washington: Executive Office of the President, October 2016) https://obamawhitehouse.archives.gov/sites/default/files/whitehouse_files/microsites/ostp/NSTC/preparing_for_the_future_of_ai.pdf; and American Bankers Association, “Understanding Artificial Intelligence” (Washington: American Bankers Association, November 2018), <https://www.aba.com/Tools/Function/Technology/Documents/understanding-artificial-intelligence.pdf>.

³ Executive Office of the President, *Preparing for the Future of Artificial Intelligence*; and Financial Stability Board, *Artificial Intelligence and Machine Learning in Financial Services* (Basel: Financial Stability Board, November 1, 2017), <http://www.fsb.org/wp-content/uploads/P011117.pdf>.

The three key components of AI--algorithms, processing power, and big data--are all increasingly accessible. Due to an early commitment to open-source principles, AI algorithms from some of the largest companies are available to even nascent startups.⁴ As for processing power, continuing innovation by public cloud providers means that with only a laptop and a credit card, it is possible to tap into some of the world's most powerful computing systems by paying only for usage time, without having to build out substantial hardware infrastructure. Vendors have made it easy to use these tools for even small businesses and non-technology firms, including in the financial sector. Public cloud companies provide access to pre-trained AI models via developer-friendly application programming interfaces or even "drop and drag" tools for creating sophisticated AI models.⁵ Most notably, the world is creating data to feed those models at an ever-increasing rate. Whereas in 2013 it was estimated that 90 percent of the world's data had been created in the prior two years, by 2016, IBM estimated that 90 percent of global data had been created in the prior year alone.⁶

The pace and ubiquity of AI innovation have surprised even experts. The best AI result on a popular image recognition challenge improved from a 26 percent error rate to 3.5 percent in just four years. That is lower than the human error rate of 5 percent.⁷ In one study, a combination AI-human approach brought the error rate down even further--to 0.5 percent.

⁴ See, e.g., Lael Brainard, "Where Do Consumers Fit in the Fintech Stack?" (speech at the FinTech Risks and Opportunities Conference, Ann Arbor, MI, November 16, 2017), <https://www.federalreserve.gov/newsevents/speech/files/brainard20171116a.pdf>.

⁵ See, e.g., Brandon Vigliarolo, "Amazon AI: The Smart Person's Guide," *TechRepublic*, August 21, 2017, <https://www.techrepublic.com/article/amazon-ai-the-smart-persons-guide/>; and Fei-Fei Li and Jia Li, "Cloud AutoML: Making AI Accessible to Every Business," *The Keyword* (Google blog), January 17, 2018, <https://www.blog.google/topics/google-cloud/cloud-automl-making-ai-accessible-every-business/>.

⁶ SINTEF, "Big Data, for Better or Worse: 90% of World's Data Generated over Last Two Years," *ScienceDaily.com*, May 22, 2013, <https://www.sciencedaily.com/releases/2013/05/130522085217.htm>; and IBM, "10 Key Marketing Trends for 2017" (2017) (on file with author).

⁷ Executive Office of the President, *Preparing for the Future of Artificial Intelligence*.

So it is no surprise that many financial services firms are devoting so much money, attention, and time to developing and using AI approaches. Broadly, there is particular interest in at least five capabilities.⁸ First, firms view AI approaches as potentially having superior ability for pattern recognition, such as identifying relationships among variables that are not intuitive or not revealed by more traditional modeling. Second, firms see potential cost efficiencies where AI approaches may be able to arrive at outcomes more cheaply with no reduction in performance. Third, AI approaches might have greater accuracy in processing because of their greater automation compared to approaches that have more human input and higher “operator error.” Fourth, firms may see better predictive power with AI compared to more traditional approaches--for instance, in improving investment performance or expanding credit access. Finally, AI approaches are better than conventional approaches at accommodating very large and less-structured data sets and processing those data more efficiently and effectively. Some machine learning approaches can be “let loose” on data sets to identify patterns or develop predictions without the need to specify a functional form ex ante.

What do those capabilities mean in terms of how we bank? The Financial Stability Board highlighted four areas where AI could impact banking.⁹ First, customer-facing uses could combine expanded consumer data sets with new algorithms to assess credit quality or price insurance policies. And chatbots could provide help and even financial advice to consumers,

⁸ AI tools are also likely to be useful for central banks and regulators in their responsibilities for supervision, financial stability, and monetary policy, although this is not addressed here. The 2017 Financial Stability Board report highlighted the potential use of AI tools by central banks and prudential authorities for applications ranging from systemic risk identification to detecting fraud and money laundering (Financial Stability Board, *Artificial Intelligence and Machine Learning*).

⁹ See Financial Stability Board, *Artificial Intelligence and Machine Learning*. See also, e.g., U.S. Department of the Treasury, *A Financial System That Creates Economic Opportunities* (Washington: U.S. Department of the Treasury, June 3, 2018), <https://home.treasury.gov/sites/default/files/2018-07/A-Financial-System-that-Creates-Economic-Opportunities---Nonbank-Financi....pdf>; Brainard, “Where Do Consumers Fit in the Fintech Stack?”; and American Bankers Association, “Understanding Artificial Intelligence.”

saving them the waiting time to speak with a live operator. Second, there is the potential for strengthening back-office operations, such as advanced models for capital optimization, model risk management, stress testing, and market impact analysis. Third, AI approaches could be applied to trading and investment strategies, from identifying new signals on price movements to using past trading behavior to anticipate a client's next order. Finally, there are likely to be AI advancements in compliance and risk mitigation by banks. AI solutions are already being used by some firms in areas like fraud detection, capital optimization, and portfolio management.

Current Regulatory and Supervisory Approaches

The potential breadth and power of these new AI applications inevitably raise questions about potential risks to bank safety and soundness, consumer protection, or the financial system.¹⁰ The question, then, is how should we approach regulation and supervision? It is incumbent on regulators to review the potential consequences of AI, including the possible risks, and take a balanced view about its use by supervised firms.

Regulation and supervision need to be thoughtfully designed so that they ensure risks are appropriately mitigated but do not stand in the way of responsible innovations that might expand access and convenience for consumers and small businesses or bring greater efficiency, risk detection, and accuracy. Likewise, it is important not to drive responsible innovation away from

¹⁰ See, e.g., Financial Stability Board, *Artificial Intelligence and Machine Learning* and National Consumers Law Center (on behalf of its low-income clients), California Reinvestment Coalition Consumer Action, Consumers Union, National Association of Consumer Advocates, U.S. PIRG, Woodstock Institute, World Privacy Forum, "Comments in Response to Request for Information Regarding Use of Alternative Data and Modeling Techniques in the Credit Process," Docket No. CFPB-2017-0005 (May 19, 2017); U.S. Department of the Treasury, *A Financial System That Creates Economic Opportunities*; and Carol A. Evans, "Keeping Fintech Fair: Thinking about Fair Lending and UDAP Risks," *Consumer Compliance Outlook* (Second Issue 2017), <https://www.consumercomplianceoutlook.org/2017/second-issue/keeping-fintech-fair-thinking-about-fair-lending-and-udap-risks/>.

supervised institutions and toward less regulated and more opaque spaces in the financial system.¹¹

Our existing regulatory and supervisory guardrails are a good place to start as we assess the appropriate approach for AI processes. The National Science and Technology Council, in an extensive study addressing regulatory activity generally, concludes that if an AI-related risk “falls within the bounds of an existing regulatory regime, . . . the policy discussion should start by considering whether the existing regulations already adequately address the risk, or whether they need to be adapted to the addition of AI.”¹² A recent report by the U.S. Department of the Treasury reaches a similar conclusion with regard to financial services.¹³

With respect to banking services, a few generally applicable laws, regulations, guidance, and supervisory approaches appear particularly relevant to the use of AI tools. First, the Federal Reserve’s “Guidance on Model Risk Management” (SR Letter 11-7) highlights the importance to safety and soundness of embedding critical analysis throughout the development, implementation, and use of models, which include complex algorithms like AI.¹⁴ It also underscores “effective challenge” of models by a “second set of eyes”--unbiased, qualified individuals separated from the model’s development, implementation, and use. It describes supervisory expectations for sound independent review of a firm’s own models to confirm they are fit for purpose and functioning as intended. If the reviewers are unable to evaluate a model in full or if they identify issues, they might recommend the model be used with greater caution or

¹¹ Lael Brainard, “Where Do Banks Fit in the Fintech Stack?” (speech at the Northwestern Kellogg Public-Private Interface Conference, April 28, 2017), <https://www.federalreserve.gov/newsevents/speech/brainard20170428a.htm>.

¹² Executive Office of the President, *Preparing for the Future of Artificial Intelligence*.

¹³ U.S. Department of the Treasury, *A Financial System That Creates Economic Opportunities*.

¹⁴ Board of Governors of the Federal Reserve System, “Guidance on Model Risk Management,” Supervision and Regulation Letter SR Letter 11-7 (April 4, 2011), <https://www.federalreserve.gov/supervisionreg/srletters/sr1107.htm>. See also Andrew Burt, “Leave A.I. Alone,” *New York Times*, January 4, 2018, <https://www.nytimes.com/2018/01/04/opinion/leave-artificial-intelligence.html>.

with compensating controls. Similarly, when our own examiners evaluate model risk, they generally begin with an evaluation of the processes firms have for developing and reviewing models, as well as the response to any shortcomings in a model or the ability to review it. Importantly, the guidance recognizes that not all aspects of a model may be fully transparent, as with proprietary vendor models, for instance. Banks can use such models, but the guidance highlights the importance of using other tools to cabin or otherwise mitigate the risk of an unexplained or opaque model. Risks may be offset by mitigating external controls like “circuit-breakers” or other mechanisms. And importantly, models should always be interpreted in context.

Second, our guidance on vendor risk management (SR 13-19/CA 13-21), along with the prudential regulators’ guidance on technology service providers, highlights considerations firms should weigh when outsourcing business functions or activities--and could be expected to apply as well to AI-based tools or services that are externally sourced.¹⁵ The vast majority of the banks that we supervise will have to rely on the expertise, data, and off-the-shelf AI tools of nonbank vendors to take advantage of AI-powered processes. Whether these tools are chatbots, anti-money-laundering/know your customer compliance products, or new credit evaluation tools, it seems likely that they would be classified as services to the bank. The vendor risk-management guidance discusses best practices for supervised firms regarding due diligence, selection, and contracting processes in selecting an outside vendor. It also describes ways that firms can provide oversight and monitoring throughout the relationship with the vendor, and considerations

¹⁵ See, e.g., *FFIEC Outsourcing Technology Services Booklet* (June 2004), <http://ithandbook.ffiec.gov/it-booklets/outsourcing-technology-services.aspx>; and Board of Governors, “Guidance on Managing Outsourcing Risk,” Supervision and Regulation Letter SR 13-19/Consumer Affairs Letter CA 13-21 (December 5, 2013), <https://www.federalreserve.gov/supervisionreg/srletters/sr1319.htm>.

about business continuity and contingencies for a firm to consider before the termination of any such relationship.

Third, it is important to emphasize that guidance has to be read in the context of the relative risk and importance of the specific use-case in question. We have long taken a risk-focused supervisory approach--the level of scrutiny should be commensurate with the potential risk posed by the approach, tool, model, or process used.¹⁶ That principle also applies generally to the attention that supervised firms devote to the different approaches they use: firms should apply more care and caution to a tool they use for major decisions or that could have a material impact on consumers, compliance, or safety and soundness.

For its part, AI is likely to present some challenges in the areas of opacity and explainability. Recognizing there are likely to be circumstances when using an AI tool is beneficial, even though it may be unexplainable or opaque, the AI tool should be subject to appropriate controls, as with any other tool or process, including how the AI tool is used in practice and not just how it is built. This is especially true for any new application that has not been fully tested in a variety of conditions. Given the large data sets involved with most AI approaches, it is vital to have controls around the various aspects of data--including data quality as well as data suitability. Just as with conventional models, problems with the input data can lead to cascading problems down the line. Accordingly, we would expect firms to apply robust analysis and prudent risk management and controls to AI tools, as they do in other areas, as well as to monitor potential changes and ongoing developments.

¹⁶ Board of Governors, "Risk-Focused Safety and Soundness Examinations and Inspections," Supervision and Regulation Letter SR 96-14 (May 24, 1996), <https://www.federalreserve.gov/boarddocs/srletters/1996/SR9614.htm>.

For example, let's take the areas of fraud prevention and cybersecurity, where supervised institutions may need their own AI tools to identify and combat outside AI-powered threats. The wide availability of AI's building blocks means that phishers and fraudsters have access to best-in-class technologies to build AI tools that are powerful and adaptable. Supervised institutions will likely need tools that are just as powerful and adaptable as the threats that they are designed to face, which likely entails some degree of opacity. While so far, most phishing attacks against consumers have relied on standard-form emails, likely due to the high cost of personalization, in the future, AI tools could be used to make internet fraud and phishing highly personalized.¹⁷ By accessing data sets with consumers' personally identifiable information and applying open-source AI tools, a phisher may be able to churn out highly targeted emails to millions of consumers at relatively low cost, containing personalized information such as their bank account number and logo, along with past transactions.¹⁸ In cases such as this, where large data sets and AI tools may be used for malevolent purposes, it may be that AI is the best tool to fight AI.

Let's turn to the related issue of the proverbial "black box"--the potential lack of explainability associated with some AI approaches. In the banking sector, it is not uncommon for there to be questions as to what level of understanding a bank should have of its vendors' models, due to the balancing of risk management, on the one hand, and protection of proprietary

¹⁷ See, e.g., Eric Lipton, David E. Sanger, and Scott Shane, "The Perfect Weapon: How Russian Cyberpower Invaded the U.S.," *New York Times*, December 13, 2016, <https://mobile.nytimes.com/2016/12/13/us/politics/russia-hack-election-dnc.html>; and Cormac Herley, "Why Do Nigerian Scammers Say They Are from Nigeria?" Microsoft Research, June 1, 2012, <https://www.microsoft.com/en-us/research/wp-content/uploads/2016/02/WhyFromNigeria.pdf>. ("A less outlandish wording that did not mention Nigeria would almost certainly gather more total responses and more viable responses, but would yield lower overall profit. Recall, that viability requires that the scammer actually extract money from the victim: those who are fooled for a while, but then figure it out, or who balk at the last hurdle are precisely the expensive false positives that the scammer must deter.")

¹⁸ Brian Krebs, "The Year Targeted Phishing Went Mainstream," *Krebs on Security*, (August 18, 2018), <https://krebsonsecurity.com/2018/08/the-year-targeted-phishing-went-mainstream/>.

information, on the other. To some degree, the opacity of AI products can be seen as an extension of this balancing. But AI can introduce additional complexity because many AI tools and models develop analysis, arrive at conclusions, or recommend decisions that may be hard to explain. For instance, some AI approaches are able to identify patterns that were previously unidentified and are intuitively quite hard to grasp. Depending on what algorithms are used, it is possible that no one, including the algorithm's creators, can easily explain why the model generated the results that it did.

The challenge of explainability can translate into a higher level of uncertainty about the suitability of an AI approach, all else equal. So how does, or even can, a firm assess the use of an approach it might not fully understand? To a large degree, this will depend on the capacity in which AI is used and the risks presented. One area where the risks may be particularly acute is the consumer space generally, and consumer lending in particular, where transparency is integral to avoiding discrimination and other unfair outcomes, as well as meeting disclosure obligations.¹⁹ Let me turn briefly to this topic.

The potential for the application of AI tools to result in new benefits to consumers is garnering a lot of attention. The opportunity to access services through innovative channels or processes can be a potent way to advance financial inclusion.²⁰ Consider, for instance, consumer credit scoring. There are longstanding and well-documented concerns that many consumers are

¹⁹ In a 2017 request for public comment, the Bureau of Consumer Financial Protection noted that alternative modeling techniques may offer consumer benefits, such as greater credit access, enhanced creditworthiness predictions, lower costs, and better service and convenience, but also highlighted consumer risks, such as privacy concerns, data quality issues, loss of the ability to correct errors, and discrimination. See Bureau of Consumer Financial Protection, "Request for Information Regarding Use of Alternative Data and Modeling Techniques in the Credit Process," (Washington: Bureau of Consumer Financial Protection, February 14, 2017), https://files.consumerfinance.gov/f/documents/20170214_cfpb_Alt-Data-RFI.pdf

²⁰ See, e.g., Lael Brainard, "FinTech and the Search for Full Stack Financial Inclusion," (speech at the Conference on FinTech, Financial Inclusion, and the Potential to Transform Financial Services at the Federal Reserve Bank of Boston, Boston, October 17, 2018), <https://www.federalreserve.gov/newsevents/speech/brainard20181017a.htm>.

burdened by material errors on their credit reports, lack sufficient credit reporting information necessary for a score, or have credit reports that are unscorable.²¹ As noted earlier, banks and other financial service providers are using AI to develop credit-scoring models that take into account factors beyond the usual metrics. There is substantial interest in the potential for those new models to allow more consumers on the margins of the current credit system to improve their credit standing, at potentially lower cost. As noted earlier, AI also has the potential to allow creditors to more accurately model and price risk, and to bring greater speed to decisions.

AI may offer new consumer benefits, but it is not immune from fair lending and other consumer protection risks, and compliance with fair lending and other consumer protection laws is important.²² Of course, it should not be assumed that AI approaches are free of bias simply because they are automated and rely less on direct human intervention. Algorithms and models reflect the goals and perspectives of those who develop them as well as the data that trains them and, as a result, AI tools can reflect or “learn” the biases of the society in which they were created. A 2016 Treasury Department report noted that while “data-driven algorithms may expedite credit assessments and reduce costs, they also carry the risk of disparate impact in credit outcomes and the potential for fair lending violations.”²³

²¹ Federal Trade Commission, “In FTC Study, Five Percent of Consumers Had Errors on Their Credit Reports That Could Result in Less Favorable Terms for Loans,” news release, (February 11, 2013), <https://www.ftc.gov/news-events/press-releases/2013/02/ftc-study-five-percent-consumers-had-errors-their-credit-reports>; Federal Trade Commission, “FTC Issues Follow-Up Study on Credit Report Accuracy,” news release, (January 21, 2015), <https://www.ftc.gov/news-events/press-releases/2015/01/ftc-issues-follow-study-credit-report-accuracy>; and Bureau of Consumer Financial Protection, *Data Point: Credit Invisibles* (Washington: Bureau of Consumer Financial Protection, May 2015), https://files.consumerfinance.gov/f/201505_cfpb_data-point-credit-invisibles.pdf.

²² The primary federal statutes governing fair lending are the Equal Credit Opportunity Act (ECOA) and the Fair Housing Act (FHA). See, also Lael Brainard, “Where Do Consumers Fit in the Fintech Stack?” (speech at the FinTech Risks and Opportunities Conference, Ann Arbor, MI, November 16, 2017), <https://www.federalreserve.gov/newsevents/speech/files/brainard20171116a.pdf>.

²³ U.S. Department of the Treasury, “Opportunities and Challenges in Online Marketplace Lending” (Washington: U.S. Department of the Treasury, May 2016), https://www.treasury.gov/connect/blog/Documents/Opportunities_and_Challenges_in_Online_Marketplace_Lending_white_paper.pdf.

A recent example illustrates the risk of unwittingly introducing bias into an AI model. It was recently reported that a large employer attempted to develop an AI hiring tool for software developers that was trained with a data set of the resumes of past successful hires, which it later abandoned. Because the pool of previously hired software developers in the training data set was overwhelmingly male, the AI developed a bias against female applicants, going so far as to exclude resumes of graduates from two women's colleges.²⁴

Importantly, the Equal Credit Opportunity Act (ECOA) and the Fair Credit Reporting Act (FCRA) include requirements for creditors to provide notice of the factors involved in taking actions that are adverse or unfavorable for the consumer.²⁵ These requirements help provide transparency in the underwriting process, promote fair lending by requiring creditors to explain why they reached their decisions, and provide consumers with actionable information to improve their credit standing. Compliance with these requirements implies finding a way to explain AI decisions. However, the opacity of some AI tools may make it challenging to explain credit decisions to consumers, which would make it harder for consumers to improve their credit score by changing their behavior. Fortunately, AI itself may play a role in the solution: The AI community is responding with important advances in developing "explainable" AI tools with a focus on expanding consumer access to credit.²⁶ I am pleased that this is one of the topics on your agenda today.

²⁴ Jeffrey Dastin, "Amazon Scraps Secret AI Recruiting Tool That Showed Bias against Women," Reuters, October 9, 2018, <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>.

²⁵ See generally ECOA section 701, 15 USC 1691; 12 CFR § 1002.9; FCRA section 615, 15 USC 1681m; and 12 CFR §§ 1022.70-75.

²⁶ See, e.g., Nanette Byrnes, "An AI-Fueled Credit Formula Might Help You Get a Loan," *MIT Technology Review* (February 14, 2017), <https://www.technologyreview.com/s/603604/an-ai-fueled-credit-formula-might-help-you-get-a-loan/>.

Looking Ahead

Perhaps one of the most important early lessons is that not all potential consequences are knowable now--firms should be continually vigilant for new issues in the rapidly evolving area of AI. Throughout the history of banking, new products and processes have been an area where problems can arise. Further, firms should not assume that AI approaches are less susceptible to problems because they are purported to be able to “learn” or less prone to human error. There are plenty of examples of AI approaches not functioning as expected--a reminder that things can go wrong. It is important for firms to recognize the possible pitfalls and employ sound controls now to prevent and mitigate possible future problems.

For our part, we are still learning how AI tools can be used in the banking sector. We welcome discussion about what use cases banks and other financial services firms are exploring with AI approaches and other innovations, and how our existing laws, regulations, guidance, and policy interests may intersect with these new approaches.²⁷ When considering financial innovation of any type, our task is to facilitate an environment in which socially beneficial, responsible innovation can progress with appropriate mitigation of risk and consistent with applicable statutes and regulations.

As with other technological advances, AI presents regulators with a responsibility to act with thoughtfulness and perspective in carrying out their mandates, learning from the experience in other areas. As we move ahead in exploring the policy and regulatory issues related to artificial intelligence, we look forward to collaborating with a broad array of stakeholders.

²⁷ See, e.g., U.S. Department of the Treasury, *A Financial System That Creates Economic Opportunities*; Brainard, “Where Do Banks Fit in the Fintech Stack?”; and Lael Brainard, “The Opportunities and Challenges of Fintech” (speech at the Conference on Financial Innovation at the Board of Governors, Washington, December 2, 2016), <https://www.federalreserve.gov/newsevents/speech/brainard20161202a.htm>.