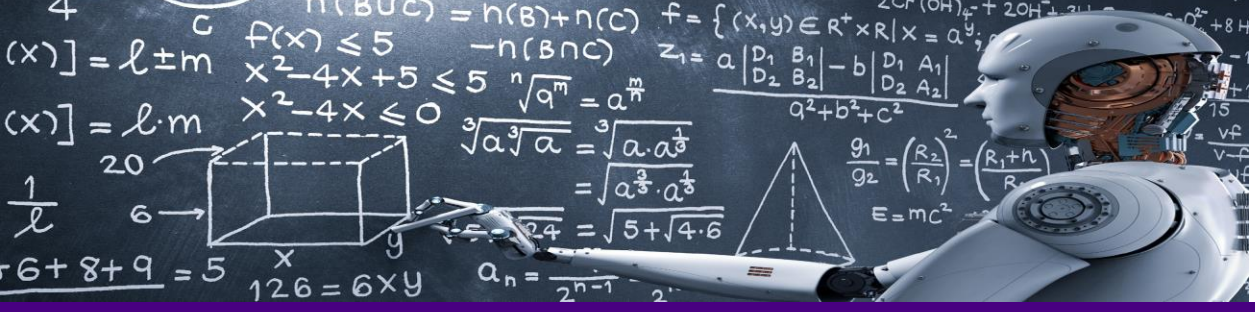

BANK POLICY INSTITUTE

Machine Learning and Consumer Banking: An Appropriate Role for Regulation

By Greg Baer and Naeha Prakash

March 14, 2019

Machine learning has the potential to democratize access to credit. It can expand the pool of people qualified to obtain credit—most notably low- and moderate-income (LMI) borrowers—and decrease the cost of that credit. It also can increase access to credit and reduce systemic risk by allowing different banks to analyze different factors, and thereby generate different results in a way that the existing, FICO-based system discourages. The greatest current obstacle to this development is pressure from the banking regulators to continue adhering to the *status quo* system, lest machine learning produce an unfortunate outcome. Perversely, that system already contains the very flaws that regulators have expressed about machine learning.

Banking regulators need to use a currently neglected tool, the notice-and-comment process required of them by Congress, to seek information and advice from experts in machine learning as to how it can benefit access to credit and what the role of regulation should be for this technology. As described below, they will likely hear that their current stance is antiquated and is the greatest current obstacle to a smart and sound way to expand credit to more Americans.

BACKGROUND

Machine learning will profoundly change how businesses make many decisions. Because machine learning works best when trained on large, disparate data sets, financial services represents an ideal environment for its application.[1] Pattern recognition algorithms can be trained to recognize not only voices and faces but also patterns of behavior, consistent with debt repayment.

While there is a recognized need for a comprehensive regulatory strategy for machine learning, innovation in banking is currently hampered by a precondition that the banking agencies "get comfortable" with machine learning in various contexts prior to its use, but particularly with respect to consumer credit.[2] We have already seen where this can lead. In the anti-money laundering context, where banks are seeking more advanced analytics to identify suspicious activity, the banking agencies have mandated that such tools be considered "models" subject to the 2011 Federal Reserve/OCC Guidance on Model Risk Management.[3] The result has been that updates to transaction monitoring programs that used to take weeks to implement now take nine months to a year.[4] Of course, that 2011 Guidance predates the revolution in

---

1 Machine learning is a subset of artificial intelligence, and connotes computer learning based on experience with a data set, with limited or no human intervention. Machine learning begins with an algorithm, which is generally a commoditized product, but its value depends crucially on the data on which it is trained.

2 For a video summary of the problem, albeit in the form of a sales pitch, you can watch https://www.zestfinance.com/zaml. It is difficult to imagine any other American industry where a vendor's presentation of an innovative product would focus primarily on how to make a compliance department and federal agency comfortable with its purchase.

3 FRB, SR 11-7: Guidance on Model Risk Management (Apr. 4, 2011); OCC, Bulletin 2011-12, Supervisory Guidance on Model Risk Management (Apr. 4, 2011).

4 Testimony of William J. Fox, House Financial Services Subcommittees on Financial Institutions and Consumer Credit and Terrorism and Illicit Finance, 9 (Nov. 29, 2017).

analytics that has occurred over the past eight years; its 21 pages of detailed mandates contain no mention whatsoever of artificial intelligence (AI), machine learning, or how examiners should review banks' use of these tools.

Nonetheless, in a recent speech, a Federal Reserve Governor stated that this Guidance should be applied to firms' use of machine learning.[5] So, too, should the Federal Reserve's Guidance on Managing Outsourcing Risk (better known as vendor management guidance), which was issued in 2013 and similarly contains no reference to AI or machine learning.[6] Indeed, machine learning algorithms are generally open source, commoditized products, and there simply is no vendor with respect to the core of the process. Fortunately, the Fed has expressed a desire not to "drive responsible innovation away from supervised institutions," and so one hopes this conversation is only beginning.

As described below, if we wish to expand access to bank credit, a very different approach is required—and quickly.

APPLICATION TO CREDIT UNDERWRITING

Summaries of machine learning are now ubiquitous, but in the financial services context, it offers an innovative method by which to improve and streamline the credit decisioning process, with the likelihood of increasing access to credit for the nearly 14.1 million underbanked individuals in the United States.[7]

Unregulated lenders, often known as fintechs, are already using AI concepts and considering a variety of factors that banks currently cannot, largely because of regulatory risk. A recent academic paper analyzes the predictive value of some of these factors by examining data from a German e-commerce company that considered both traditional credit scores and personal information gleaned from the customer's digital footprint.[8] The digital footprint included seemingly tangential factors, such as whether the borrower orders products from a desktop or a phone; if a phone, which type of phone; the time of day of a purchase; and whether a purchaser comes to a site from a price-comparison website or from a linked advertisement. These factors proved highly probative; for example, customers whose names are contained in their email addresses are 30% less likely to default than those who do not.

Indeed, and startlingly, that paper finds that the digital footprint variables it examined were collectively more accurate than credit scores in predicting consumer default. Furthermore, while many LMI borrowers may not have the credit history to even obtain a credit score, they are more likely to have a digital footprint. Finally, unlike a credit score, this data is practically costless to obtain for any lender, which should translate to lower rates on loans. Thus, the paper finds that "digital footprints can facilitate access to credit when credit bureau scores do not exist, thereby fostering financial inclusion and lowering inequality."[9]

SOME PERSPECTIVE

The fact that we are even having this debate with respect to bank credit is somewhat remarkable, given the regulatory status quo. Currently, for a consumer looking to obtain a loan, a key determinant is a formula developed in the 1980s by Fair Isaac & Company, a data analytics company in San Jose, California, commonly referred to as the FICO score. (The same is true for one seeking to rent an apartment, obtain insurance, or get a job.) Each of the three large, national credit-reporting agencies has developed proprietary variants of FICO scores, which are used to inform at least 90 percent of the lending decisions made in the United States."[10]

---

5 Brainard, L. *What Are We Learning about Artificial Intelligence in Financial Services*? (Nov. 13, 2018).

6 FRB, Supervisory Letter 13-19/ CA 13-21: Guidance on Managing Outsourcing Risk, (Dec. 5, 2013).

7 FDIC, National Survey of Unbanked and Underbanked Households (2017).

8 Berg, Burg, Gombovic, and Puri, *On the Rise of the FinTechs—Credit Scoring using Digital Footprints* (Sept. 2018).

9 *Id.* at 6.

10 Report to Congress on Credit Scoring and Its Effects on the Availability and Affordability of Credit (Aug. 2007).

Although the credit reporting agencies have disclosed the basic inputs (*e.g.*, payment history, debt burden, length of credit history, types of credit used, and recent searches) into the FICO model, it is unknown how those components are combined or weighted.[11] While the CFPB has knocked on the door of the consumer reporting agencies to assess overall regulatory compliance,[12] neither it nor the Department of Justice appear to have taken steps to subpoena credit score formulas to assess their alignment with anti-discrimination laws or other Federal consumer financial protection laws. For their part, the banking regulators have not, to our knowledge, sought to determine whether the formula makes them comfortable from a safety and soundness perspective. They certainly have not required banks to incorporate FICO algorithms as part of their model review process, as mandated under their Guidance on Model Risk Management. Nonetheless, the U.S. financial regulators continue to permit FICO's use in extending credit to consumers, and the banking regulators have even explicitly incorporated these scores into their capital requirements.

Rather than insisting on monitoring the *process* for developing FICO-based algorithms, *ex ante,* the banking agencies have evaluated its *results, ex post*.[13] Thus, banks are responsible for examining their lending patterns to identify not only poor underwriting but also potential violations of fair lending laws; banking regulators examine and validate that work through the examination process.

This history is portable to AI and machine learning. Indeed, the CFPB took a similar approach when assessing Upstart for a No-Action Letter on its automated underwriting model.[14] In that example, the CFPB reviewed Upstart's automated underwriting model, including the underlying alternative data used, *after* its design and use, and determined that the model did not warrant, at present, enforcement or supervisory action under the Equal Credit Opportunity Act.[15]

WHY MACHINE LEARNING AND CONSUMER CREDIT MATTER

The stakes involved with how regulators treat AI and machine learning are extremely high. Machine learning has the potential to cure two great flaws in the existing, FICO-based system: first, the FICO-based system denies bank credit to a significant portion of the population, generally LMI borrowers, who are underserved or underbanked and therefore unable to generate a FICO score. Banks, which are the low-cost provider of credit, generally rely primarily on FICO scores to offer consumer credit, as they come with effective immunity from litigation or enforcement on fair lending grounds. Second, the FICO-based system reinforces the same bias in every bank, which not only constrains credit, but also creates concentration risk, and therefore systemic risk.

If regulators follow the same approach they have with FICO, machine learning has the potential to vastly expand access to credit.

First, as noted above, machine learning allows credit decisions to be made based on a wide range factors beyond a consumer's prior credit history, thereby expanding the population eligible to receive bank credit. Under the existing framework, certain minority communities, whether based on their race, ethnicity or geography, have been

---

11 Consumers have some rights with respect to their FICO score under the Fair Credit Reporting Act—namely, the right to dispute inaccuracies in the data that is input into the formula—but they have no right to see the formula itself. GAO, Credit Reporting Literacy (Mar. 2005).

12 CFPB, Consumer Reporting Examination Procedures (Sept. 2012).

13 A 2007 Board report to Congress states: "In the broadest terms, the findings of the Federal Reserve study are as follows: (1) The credit history scores evaluated here are predictive of credit risk for the population as a whole and for all major demographic groups. That is, over any credit- score range, the higher (better) the credit score, the lower the observed incidence of default. These conclusions are limited to credit history scores, that is, scores calculated exclusively on the basis of individuals' credit records as assembled by the three national credit-reporting agencies (Equifax, Experian, and TransUnion). Other kinds of credit scores were not studied here."

14 CFPB, Upstart Request for No-Action Letter (2017).

15 CFPB, Upstart No-Action Letter (2017).

disproportionately left behind in the credit scoring process.[16] For example, a low-income person may not have sufficient credit history to merit a good FICO score, and a recent immigrant receives no credit for any credit history outside the United States. But other factors, such as the applicant's current employment or educational information, would allow for more accuracy regarding that individual's financial ability or repayment propensity.[17] In addition to reaching consumers with little to no credit history, machine learning could produce benefits for individuals with previously damaged credit, who would benefit from credit scores incorporating information related to, for example, the use of credit counseling services.

Machine learning would not only expand the pool of borrowers eligible for bank credit but also would lower its cost, thereby expanding the number of people able to access these services. As the Financial Stability Board described in a 2017 report, as banks turn to alternative data sources, "[a]pplying machine learning algorithms to this constellation of new data has enabled assessment of qualitative factors such as consumption behaviour and willingness to pay. The ability to leverage additional data on such measures allows for greater, faster, and cheaper segmentation of borrower quality and ultimately leads to a quicker credit decision."[18] Streamlining the credit decisioning process can allow for these costs to be lowered and translated to new, innovative consumer financial products and services, which serve a broader range of borrowers.

Second, machine learning would diversify views of credit risk across banks, both based on the bank's lending profile and the consumer's experience, and would be driven by each bank's underwriting practices and risk management policies. What is remarkable about a FICO score is not only its opacity but its ubiquity. A bad FICO score reduces a consumer's chances of obtaining credit not just from one bank, but from *all* banks. This not only reduces the availability of credit to potentially creditworthy borrowers, but also tends to homogenize bank balance sheets around those who do score well. With machine learning, each bank can train its algorithm on its own data sets and adjust the algorithm according to its own risk tolerance and risk management practices. A FICO score will no longer be destiny, and banks will take different views of credit.

Machine learning also allows for a more dynamic view of credit than a FICO score, thereby expanding an existing customer's ability to obtain future credit. In particular, machine learning offers a real-time process to factor in a broader set of attributes and experiences to predict certain outcomes, including default, delinquencies or repayment.[19] While credit reporting agencies no doubt conduct ongoing assessments to determine the predictability of its scores, and regularly adjust the variables to produce more accurate results, a credit score is ultimately a static formula.[20] Machine learning, alternatively, is adaptive and allows for the discovery and use of more factors, resulting in greater accuracy regarding an individual's credit risk.

THE REGULATORY APPROACH

None of the above is to suggest that machine learning is a panacea and that there is not a potential for bias. By now, everyone knows the anecdote of Amazon using machine learning for its hiring decisions, where the algorithm advised

---

16 CFPB, Data Point: Credit Invisibles (May 2015).

17 CFPB, Upstart Request for No-Action Letter (2017).

18 FSB, Artificial Intelligence and Machine Learning in Financial Services: Market Developments and Financial Stability Implications (Nov. 1, 2017).

19 Press, G, *Equifax and SAS Leverage AI and Deep Learning to Improve Consumer Access to Credit*, FORBES (Feb. 20, 2017).

20 Fair Isaac has recently released an "UltraFICO" score, which will supplement its existing FICO score if a consumer falls short. As described by Fair Isaac, several factors could improve a consumer's UltraFICO score: having an account open for a long period of time; having a positive cashflow; paying bills through checking accounts; having at least several hundred dollars in their accounts.

hiring only men. Machine learning can identify spurious correlations and begin making recommendations based on them.[21] It also can identify invidious correlations and begin making recommendations based on them.

As one practitioner explains, "[T]he benefit of a computer is that it can find non-obvious correlations between patterns of human behavior we may not have considered and the behavior we are concerned with, like defaulting on a loan, which can make our businesses more efficient. It can also be terribly, terribly wrong.... For example, a computer could find correlations between a person's race and his or her likelihood to launder money. The computer would, without guilt or hesitation, make transaction recommendations based on these metrics. But these recommendations would not only be unethical, they would factually mislead the human decision maker." [22]

The answer is simple, and precedented. Assisted AI or machine learning incorporates humans in the process of analyzing the results that machine learning produces, through back-testing, trend analysis and exception reviews, correcting any spurious or invidious results. Presumably, this process is not significantly different from how banks ensure that their FICO models are not producing unexpectedly high defaults or unforeseen disparate impact outcomes.

Widespread use of machine learning also would have a self-correcting feature that the current, FICO-dependent model does not. If one bank were to begin making poor decisions, others would remain available to consumers and businesses, because of greater diversity in credit views.

When Amazon had its now infamous glitch, no one suggested as the remedy for the Department of Labor to review all future Amazon's human resources technologies. Rather, Amazon changed the program when it saw these results. If it had failed to do so, presumably one of many possible governmental agencies would have taken action to force it to do so. Likewise, a bank should be empowered to develop and manage the risks related to its use of AI or machine learning, and change their programs where warranted or based on feedback from its regulators, as is current practice.[23] Similar to current fair lending protocol, bank examiners would monitor that process.

In short, there should be no need for regulatory agencies to issue a No-Action letter, engage in a "sandbox" or "innovation strategy," or require examiner approval for every AI or machine learning application. Rather, they simply need to examine credit underwriting using these concepts the way they have with the FICO system—by examining banks as they back-test and validate results.

A SIMPLE FIRST STEP

As with most regulatory policies, outcomes here could be much improved through transparency and adherence to the rule of law. The Federal Reserve/OCC Model Risk Management Guidance runs to 21 pages of detailed prescriptions, and is enforced as a binding regulation through the examination process; *it is the greatest current obstacle to the banking industry using AI and machine learning to expand access to credit.* Nonetheless, it was neither issued for public comment under the Administrative Procedure Act nor submitted for Congressional review under the Congressional Review Act. Even if the guidance made sense for the models in use in 2011—and there is good reason

---

21 For some amusing examples, or if you are a Nicholas Cage fan.

22 Shiffman, G, *The Challenge of Artificial Intelligence* (Feb. 19, 2019).

23 Governor Brainard's speech on AI in financial services suggests that risk-based approach should be taken with respect to the use of these tools, which seemingly acknowledges that banks are best placed to understand and manage the risks involved with using AI or machine learning.

to believe that it did not—it now represents *exactly the wrong way to think about AI and machine learning*. The same is generally true with respect to the agencies' vendor management guidance.[24]

Thus, a good first step in deciding how regulation and examination should be adapted to the world of machine learning would be for the banking agencies to seek comment, through an advanced notice of proposed rulemaking, on how they should examine banks' use of AI and machine learning. There is a wealth of information available in this area, and a multiplicity of experts whose insights would no doubt be interesting and valuable—if the agencies would simply seek it out. Since 1946, Congress has effectively required federal agencies to engage in open-source rulemaking, and conformance to that law for regulation of AI and machine learning would be wise, even poetic.

Disclaimer: The views expressed in this post are those of the author(s) and do not necessarily reflect the position of the Bank Policy Institute or its membership, and are not intended to be, and should not be construed as, legal advice of any kind.

[24] The Federal Reserve's guidance, described above, is a model of brevity when compared to the hundreds of pages of guidance issued by the Office of the Comptroller of the Currency. For a bibliography of all the relevant bulletins, guidance, and advisory letters, see OCC Bulletin 2013-29. As with the Federal Reserve, the OCC's various publications on vendor management, although articulated, especially recently, as guidance, can often be treated as binding rules in the examination process. Publishing these for public comment would be a benefit to consumers. (It would also be a benefit to the numerous small businesses who have been effectively foreclosed from serving banks, as they cannot meet the immense compliance burdens that come with doing so, and banks have every incentive to minimize the number of vendors they need to run through the regulatory gauntlet.)