

Assessing Federated Machine Learning's Potential for Transforming KYC/AML *A Proposal to the Central Bank of the Future Project*

Introduction

This proposal describes an interdisciplinary research approach to assess the potential of using federated machine learning models to both improve the effectiveness of anti-financial crimes (AFC) risk management and to enhance the inclusiveness of the global financial system.¹ Federated machine learning has the potential to help firms improve the accuracy and predictive power of models that identify suspicious transactions and make risk management decisions about whether and how to do business in certain jurisdictions on the basis of better information about actual financial crimes risks. The impact of this privacy-enhancing technology may be particularly transformative for countries and individuals that have lost access to the global financial system due to de-risking decisions by individual firms and diminished correspondent banking. Our approach to assessing this potential is to demonstrate a prototype of federated learning for AFC² while engaging regulators, industry, and advocates to understand the foundational principles, specific requirements, and key policy debates that this emerging technology needs to serve. By beginning this engagement at the start of our research, we ensure the creation of actionable evidence that speaks to the needs of all stakeholders and has the potential to drive widespread adoption and responsible use.

In a federated learning approach, an AML model might be trained separately at each financial institution, but model parameters³ can be shared via a central authority to improve overall performance. This approach is particularly promising precisely because it enables the benefits of machine learning – better pattern identification and enhanced predictive power – without requiring more data sharing or compromising compliance with privacy expectations, data localization laws, or law enforcement confidentiality considerations.

Federated machine learning can leverage data insights currently beyond the scope of many individual firms' AFC programs, let institutions learn from each other without sharing sensitive information, and enable law enforcement and regulatory agencies to safely provide feedback on firms' risk identification

¹ AFC as used in this proposal reflects the umbrella term increasingly used by practitioners around the world to refer to activities to implement requirements of the Bank Secrecy Act and similar laws in other jurisdictions. These include requirements related to anti-money laundering (AML), countering the financing of terrorism (CFT), and sanctions monitoring, as well as customer identification or know-your-customer (KYC) requirements. KYC programs typically include customer identification programs (CIP) to conduct customer due diligence (CDD) and enhanced due diligence (EDD) as needed.

² Several technology companies have developed approaches using federated machine learning to help financial services firms with respect to AFC expectations. The research proposed herein could be structured to evaluate their products or work with one of those companies to develop a prototype for evaluation in transaction monitoring or related AFC disciplines.

³ A model parameter is a configuration variable that is internal to the model and whose value can be estimated from data. They are the part of the model that is learned from historical training data.

efforts. Notably, federated learning can accomplish this without requiring direct sharing of data among competing firms or by government agencies.

Consideration of federated machine learning in the anti-financial crimes context also focuses attention on a critical role that central banks around the world can play in fostering a more inclusive global financial system. In this regard, central banks in developing countries, acting individually or in groups, may have a particularly compelling opportunity beyond writing rules to help speed adoption of new technologies that facilitate more accurate and inclusive screening of their activities within their borders. Indeed, in federated learning such institutions could serve as a utility-like hub that aggregates and shares model parameters and other derived insights to regulated entities. Taking on this role can give central banks a pivotal lever to achieve two important goals:

- Improving domestic access to the international financial system, especially in areas critical for growth and development like foreign direct investment, remittance transfers, and other money services businesses
- Enhancing confidence on the part of investors, firms, and international bodies in the ability of domestic actors to participate safely and responsibly in the global financial system.

For this project, we propose to evaluate federated machine learning technology designed to detect a range of illicit activities, including money-laundering activities. This research will:

- Evaluate the performance of federated learning technology in the context of financial crime risks (including how well it reduces improper denials of access that result from false positives);
- Assess the performance of federated machine learning models in reliably protecting sensitive information while enabling sharing of model parameters and other data insights;
- Measure how the use of federated machine learning affects individuals and businesses underserved by current approaches (using metrics such as income, geography, and other variables); and
- Identify policy options for central banks and other financial regulators to support the use of technologies that improve financial inclusion.

We begin with a brief overview of the market and regulatory context relevant to the use of federated machine learning in AFC risk detection before providing an overview of our proposal for research in this area.

Market Context

The Problem of Financial Exclusion

Several factors have combined to exclude significant populations, disproportionately in emerging markets, from the global financial system. Two stand out as key gating factors:

- The inherent difficulty and additional cost of risk-rating activity that occurs outside mainstream forms of finance, especially in connection with money services businesses
- The regulatory risks and reputational impact for financial institutions related to being connected to illicit financing activities.



Exiting markets and lines of business to de-risk operations has occurred extensively over the past decade. This has disproportionately affected countries and regions where data limitations and existing market structures limit firms' ability to identify individuals and businesses for bank access to the required standard or increase the costs and risks related to doing so. This in part means that approximately 30% of the adult population globally does not have an account from which they can safely transact, save, or access credit.⁴ As a result, almost two billion adults around the world are unbanked, and are likely to rely on products that pose greater financial risks.⁵

Impact on Developing Countries

In the context of developing economies, technology is increasingly a focal point for considering how policy can accelerate development and convergence to international business and risk management norms. The former President of the Financial Action Task Force (FATF) recently recognized the important role that technology can play in improving global AFC efforts:

[F]intech can be a significant contribution to a more effective anti-money laundering and counter-terrorism-finance policy. If we strike the balance right, and we engage the industry and the industry designs its products correctly that takes it into account, we will have a much more effective policy against terrorism financing.⁶

As currently implemented, firms often cite compliance costs related to AML and CFT regulations and the potential for sanctions for noncompliance as key factors in decisions about whether and how they serve developing countries.⁷ This is particularly acute for firms providing correspondent banking services in emerging markets where data availability and systems limitations make AFC processes more time consuming and expensive, as well as potentially less effective in identifying illicit activity.⁸ For example, correspondent banking relationships (CBR) have shrunk globally by 20 percent in the last seven years.⁹ One survey found that 21 Angolan banks faced relationship termination with foreign banks due to cited CDD compliance costs and challenges related to money laundering risks.¹⁰ Another survey showed that 21

⁴ Asli Demirgüç-Kunt, Leora Klapper, Dorothe Singer, Saniya Ansar, and Jake Hess, Global Findex Database 2017: Measuring Financial Inclusion and the Fintech Revolution, World Bank Group (2018). Traditional estimates of individuals who lack access to mainstream financial products and services may understate the effect of current policies and regulations on financial inclusion and asset building. Such estimates do not typically capture the dampening effect customer screening and other risk management processes can have on correspondent and commercial banking activities which are critical to enabling broader economic development and growth.

⁵ *Ibid.*

⁶ Samuel Rubinfeld, "FATF Wants to Use FinTech Against Money Laundering," The Wall Street Journal, June 12, 2017, <https://blogs.wsj.com/riskandcompliance/2017/06/12/fatf-wants-to-use-fintech-against-money-laundering/>

⁷ "The Decline in Access to Correspondent Banking Services in Emerging Markets: Trends, Impacts, and Solutions," The World Bank, 2018, <http://pubdocs.worldbank.org/en/786671524166274491/TheDeclineReportlow.pdf>

⁸ "The Decline in Access to Correspondent Banking Services in Emerging Markets: Trends, Impacts, and Solutions," The World Bank, 2018, <http://pubdocs.worldbank.org/en/786671524166274491/TheDeclineReportlow.pdf>

⁹ Shirish Wadivkar, "De-Risking in Correspondent Banking: Existential Challenge or Catalyst for Change?" Standard Chartered, 2019, <https://www.sc.com/en/feature/de-risking-in-correspondent-banking-existential-challenge-or-catalyst-for-change/>

¹⁰ "The Decline in Access to Correspondent Banking Services in Emerging Markets: Trends, Impacts, and Solutions," The World Bank, 2018, <http://pubdocs.worldbank.org/en/786671524166274491/TheDeclineReportlow.pdf>



of 23 Caribbean banks in 12 countries had lost at least one CBR for similar reasons.¹¹ Reduction in correspondent banking activity has a particularly acute effect on financial inclusion, because these activities include remittance and foreign direct investment.¹² Restricting these activities can stunt development, growth, and asset formation, as well as cutting off individuals from immediately accessible financial resources from family members living abroad. It may also push people and businesses to use less transparent and unregulated financial networks.¹³

As discussed more completely below, federated machine learning has the potential to improve the interplay of AFC requirements and the needs of emerging markets. Firms in several developing countries are already piloting machine learning-based technologies for AML and KYC processes. For example, one UK bank has begun deploying ML technology for AML across Africa.¹⁴

The Potential of Federated Machine Learning

Emerging technologies, however, present significant potential to expand access to the financial system through a set of interrelated improvements in AFC risk management processes:

- Enhancing the accuracy of the processes used to evaluate individuals and businesses at the application stage and to detect illicit transactions
- Improve the ratio of true and false positives, which will focus investigative resources on real risks and may reducing the costs that firms incur related to customer screening and risk monitoring;
- Reducing obstacles to collaboration among competing firms
- Enabling firms to enter and serve new markets responsibly.

Machine learning and other forms of artificial intelligence can enable firms and regulators to leverage data insights across financial firms, to expand the kinds of data used for customer screening and risk identification, to identify patterns and activities that are not currently detectable, and to adapt to changing data and market conditions more swiftly than with current methods.

Federated machine learning may be the technology with greatest promise to enable these kinds of advances in AFC risk management. This form of artificial intelligence can enable sharing insights from data under long-established, but rarely used, provisions of the Patriot Act and enhance the accuracy and efficiency of systemic efforts to detect and prevent illicit financial activity.¹⁵ Notably, these benefits can be obtained without requiring firms or government agencies to share sensitive information. This approach

¹¹ Scott B. MacDonald, "Is There a 'New Normal' for De-Risking in the Caribbean," The Center for Strategic and International Studies, October 16, 2019, <https://www.csis.org/analysis/there-new-normal-de-risking-caribbean>

¹² Amounting to around \$800 billion for lower-middle income countries in 2018. Given that remittances and foreign direct investment are about five times higher than official development assistance (around \$150 billion in 2018), the development impact of restricting cash transfers when CBRs are severed could be drastic. Dilip Ratha, "Remittances on Track to Become the Largest Source of External Financing in Developing Countries," World Bank Blogs, April 8, 2019, <https://blogs.worldbank.org/peoplemove/remittances-track-become-largest-source-external-financing-developing-countries>; "Development Aid Drops in 2018, Especially to Neediest Countries," OECD, October 2019, <https://www.oecd.org/development/development-aid-drops-in-2018-especially-to-neediest-countries.htm>

¹³ *Ibid.*

¹⁴ Artificial Intelligence for Africa: An Opportunity for Growth, Development, and Democratization," Access Partnership, March 2019, https://pic.strathmore.edu/wp-content/uploads/2019/03/PIC_AI_for_Africa_Whitepaper.pdf

¹⁵ 31 U.S.C. 5311.



also has promise for enabling compliance with emerging privacy frameworks and data localization laws that are beginning to create barriers to data flowing across borders. Federated learning can be used to improve the risk management processes that firms use throughout the AFC lifecycle by allowing firms to use model parameters that reflect the experience of peers and government agencies.

Federated machine learning might improve transaction monitoring so that firms can identify true risks in greater volume and earlier, while reducing false positives. These improvements may, in turn, improve client screening models, as well as fraud models, by providing more refined and accurate information about who is likely to engage in illicit financial activity. Enabling better risk detection throughout the AFC lifecycle has the potential to make each individual process more efficient, more cost effective, and more inclusive.

But there remains significant research to be done on how this technology can be applied in the financial crimes context, including evaluating whether and how it can be designed to be trusted to reliably identify candidates that should be permitted to transact in the global financial system; developing evidence to support improvements in financial inclusion; and inform the evolution of policy to support enable responsible innovation.

Regulatory Requirements

Since terrorist attacks on the United States in 2001, the global legal and regulatory infrastructure for preventing illicit financial activity has been greatly strengthened. Although governments and banking regulators around the world have played a role in articulating new requirements, a small number of countries, with the United States foremost among them, have disproportionately shaped global standards given the interconnected nature of financial transactions. In some cases, U.S. or European standards are effectively exported because some aspect of a transaction touches those jurisdictions or other countries have harmonized their law and regulation to those standards to promote investor confidence in their markets.

In Section 314(b),¹⁶ the Patriot Act also recognized the systemic importance of sharing information among firms to effective prevention of financial crime. That section enables financial institutions to share information among themselves in the case of suspected money laundering or terror financing. Given that notification is provided to FinCEN, these acts of sharing customer data are generally allowed under a safe harbor from liability of violating other laws including the Right to Financial Privacy Act and the Gramm-Leach-Bliley Act. However, in practice little information sharing occurs under this provision due primarily to concerns related to privacy and competition, as well as lack of clarity in interpreting the necessary legal predicate.

The Bank Secrecy Act generally requires banks in the U.S. to file a Suspicious Activity Report (SAR), when there is suspicion of that an activity has no business purpose or no apparent lawful purpose.¹⁷ In addition, Section 314(a) of the USA PATRIOT Act of 2001 (Patriot Act) and its implementing regulations require financial institutions to provide details on accounts and transactions when law enforcement requests that

¹⁶ *Ibid.*

¹⁷ A large number of countries around the world have similar activity reporting requirements focused on a domestic regulatory or law enforcement agency.



data for suspected money launderers or terror financiers.¹⁸ However, firms making these reports receive little to no feedback on the accuracy of risks identified in SARs,¹⁹ which limits their ability to fine tune models. Further, supervision procedures can in some respects incentivize firms to file high volumes of SARs regardless of the accuracy of their reports.

Against this backdrop, the creation of privacy frameworks for the digital age can present new obstacles for the kind of data sharing that the Patriot Act was designed to encourage.²⁰ For example, individual countries are developing standards that define when and how their citizens can control the use of their digital data and may require that the individual be given notice of and consent to a specific use of their data. In the same vein and often with the intent of reinforcing national privacy rules, data localization or residency laws require that data about a nation's citizens or residents be collected, processed, and/or stored inside the country, and may require demonstration that local requirements have been satisfied before data can be transferred internationally. India, China, and Russia have moved in this direction, as have many African nations.

Developing Countries Context

Although the Bank Secrecy Act and similar U.S. laws are broadly influential, every country maintains slightly different AFC regulatory frameworks promulgated by their governments or central banks. In trying to navigate these differences when managing cross-border transactions, firms have often adopted compliance frameworks that implement a universal set of minimum requirements that reflect the most restrictive approach applicable to any jurisdiction in which they operate. This may result in unduly conservative positions on specific interpretative issues. Similarly, mixed signals about regulatory interpretations and priorities have "result[ed] in simplistic risk assessment methodologies being applied by these entities"²¹ and wholesale decisions to exit product markets and regional markets.

An Approach to Federated Machine Learning For Financial Crimes

Early evaluations of federated machine learning technology, especially in the area of screening applications and transactions for fraud, suggest it can increase the efficiency, effectiveness, and fairness of risk detection processes. Specifically, this technology-driven approach to illicit activity detection not only improves true positive rates, but also potentially improves consumer data privacy, supports adherence to country-specific data localization laws, and reduces improper exclusion of particular populations.²²

¹⁸ 31 U.S.C. 5311; *see also* 31 C.F.R. Part 1010.520.

¹⁹ Joe Mont, "Data Sharing, AI May be Antidote to Failing AML Efforts," Compliance Week, Jan 17, 2018. <https://www.complianceweek.com/data-sharing-ai-may-be-antidote-to-failing-aml-efforts/2411.article>

²⁰ Jonah Anderson, Jeremy Kuester, John Wagner, Rebecca Copcutt, and John Timmons, AML Information Sharing in a Technology-Enabled and Privacy-Conscious World, White & Case, (January 31, 2019).

²¹ Clay Lowery and Vijaya Ramachandran, "Unintended Consequences of Anti-Money Laundering Policies for Poor Countries," Center for Global Development, November 9, 2015, <https://www.cgdev.org/publication/unintended-consequences-anti-money-laundering-policies-poor-countries>

²² Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith, Federated Learning: Challenges, Methods, and Future Directions, arXiv preprint arXiv:1908.07873, (Aug. 2019); Toyotaro Suzumura, Yi Zhou, Nathalie Baracaldo, Guangnan Ye, Keith Houck, Ryo Kawahara, Ali Anwar, Lucia Larise Stavarache, Yuji Watanabe, Pablo Loyola, Daniel Klyashtorny, Heiko Ludwig, and Kumar Bhaskaran, Towards Federated Graph Learning for Collaborative Financial



How Federated Machine Learning Works

In general, federated learning refers to a system that shares pattern information on machine learning models from different financial institutions in such a way that it enables tuning of the associated model parameters across financial institutions. The distinguishing characteristic of federated learning is that the multiple different machine learning models are each trained “locally” on segregated information by individual firms using their own customer data. Once trained, individual firms share details about the resulting model parameters with the other models in the federation of institutions, but not the underlying data from which those parameters were derived.²³ Typically, the federated models use a common machine learning algorithm and common input variables so that a master model can be created by combining the model weights through averaging, for example. However, as may be the case with AML models developed independently at different firms, the models may be built using different algorithms and inputs. In this case, the different models can be combined to create a master model that is an ensemble of distinct models, as opposed to the standard federated approach where the master model is a single model created from combining internal parameters.

The key benefit of federated learning is that it effectively expands the data set on which the detection algorithm is trained without actually exchanging the underlying data. In AFC, a machine learning model is typically trained on a large volume of data which includes both transaction data that form the inputs to the detection system as well as “truth” data, which indicates, for each transaction, whether it corresponds to illicit or benign activity. The machine learning model adapts its internal parameters, or weights, to optimize its ability to classify a set of inputs as either benign or illicit. Such machine learning techniques require vast amounts of data showing both positive cases (in this context would be actual illicit activity) and negative ones (here, benign transactions). In practice, firms have limited access to training data about illicit activity because they are unlikely to experience large volumes of illicit activity directly and because governments provide little, if any, feedback regarding accuracy or validity of SARs to individual firms or the industry as a whole.

How Federated Machine Learning Can Improve Financial Inclusion

There are several ways in which federated machine learning can be employed to improve global financial inclusion. There is the basic benefit of reduced cost and uncertainty of having automated models with a central bank closely involved in standards and operations. Financial institutions in less developed countries or regions can also benefit from patterns that have been learned from regions with the resources to have more mature and extensive AML processes, thus effectively jump-starting the AML process in the developing country. In return, this process can provide the developed countries more confidence in the AML processes of the developing regions precisely because they are building off of trusted models from the developed countries. That greater trust can then engender increased financial investment in the developing country and accelerate development.

Beyond enabling more data for model training, federated machine learning can potentially provide valuable insights into classifiers for particular populations. For example, if one financial institution has a

Crimes Detection, IBM TJ Watson Research Center, (2019) (discussing work at the United Kingdom Financial Conduct Authority week-long Global Anti-Money Laundering and Financial Crime TechSprint done by interdisciplinary teams to develop preliminary proofs-of-concept of several technologies for, among other things, applications of federated machine learning to assess various money laundering risks and to validate beneficial ownership).

²³ Qiang Yang, Yang Liu, Tianjian Chen, and Yongxin Tong, Federated Machine Learning: Concept and Applications. ACM Transactions on Intelligent Systems and Technology, 10:12, (January 2019).



majority of customers from a particular disadvantaged population or region, the results of that institution's local AML detection will likely determine the best distinguishing features to accurately identify good and bad actors within that population. A different financial institution that has a much smaller proportion of customers from the same disadvantaged population or region will likely not have enough data on that population to determine specific distinguishing features and may result in flagging the entire population as high risk, even if it would otherwise like to do business there. Sharing model parameters calculated from the financial institution with that population as a majority can improve the second institution's performance in screening that population for risks of illicit financial activity and reduce barriers to entry in markets that are underserved.

This federated approach avoids explicitly sharing customer data while capturing the distinguishing aspects of the data distributions to distinguish between benign and illicit activity. However, the technology is still maturing, and there are risks that need to be investigated and managed. For one, such systems must be carefully designed and tested to ensure private or sensitive information is not inadvertently leaked in the process of sharing model updates. This form of aggregating the data statistics appears, on the surface, to preserve privacy in that it seems unlikely that parameters based on model that have been trained on large volumes of data could expose details of an individual's data, but a central bank, or other regulator should be aware of the potential for leaking sensitive information.²⁴ Another concern relates to inaccuracies in shared model parameters that could poison the master model as has been explored in recent research evaluating such an attack on a federated learning system.²⁵

Rethinking the Role of the Central Bank of the Future

Preliminary conversations with global regulators and industry stakeholders have identified the need to understand how novel technologies for financial institutions, regulators, and law enforcement can help these stakeholders collaborate more efficiently to improve AFC processes. Just as compelling is the need to overcome significant obstacles to such collaboration. These obstacles result from conflicts and ambiguities in current policy, from privacy and competitive concerns, and from uncertainty about the reliability and performance of new approaches. In the context of AFC risk identification, the latter presents a particularly acute form of the innovator's dilemma in developing a foothold from which to scale, since firms whose AFC processes fail may face financial losses, regulatory sanction, and reputational damage. Indeed, the standard of proof in the context of financial crimes context that must be satisfied before a firm makes a significant investment in federated learning or other innovation is likely to be higher than in other applications, like fraud detection. Given the risks for firms in adopting a technology that has not been fully vetted or proven at scale, a central bank can make a significant contribution to accelerating progress in this area by establishing a rigorous basis for understanding how such approaches work and what the benefits and costs of using them are.

²⁴ In at least one documented situation, researchers were able to recover private information (in this case a credit card number) from the parameters shared among the federated learning models in the context of a specific learning algorithm that sought to predict the next character typed on a smartphone. See Nicholas Carlini, Chang Liu, Úlfar Erlingsson, Jernej Kos, and Dawn Song, *The Secret Sharer: Measuring Unintended Neural Network Memorization and Extracting Secrets* (2018).

²⁵ Minghong Fang, Xiaoyu Cao, Jinyuan Jia, Neil Zhenqiang Gong, *Local Model Poisoning Attacks to Byzantine-Robust Federated Learning*, *Usenix Security Symposium*, (Feb. 2019).



Furthermore, the development of an AFC ecosystem that leverages federated learning may point to a utility-like role for a central bank or set of regional central banks. For example, a central bank could act as an aggregator of models and/or model parameters, where each financial institution periodically shares details of its models (explicit weights, gradients, or other updates) with the central bank (or other central authority). The central bank would then combine the shared model details to create the master model and distribute the resulting master model updates back out to the participating financial institutions. Each separate model can then be updated according to its own data insights in addition to the parameters shared by the central bank.²⁶ For the central bank of a developing country or a group of such institutions, providing this kind of support has the potential to accelerate convergence of AFC processes among supervised firms to international norms and to foster trust among firms that may otherwise hesitate to do business in the region and before international bodies such as FATF.

Proposed Evaluation of Federated Machine Learning

Through conversations with industry and regulatory stakeholders, we have identified that improving information sharing through the application of new technologies for AFC risk detection could significantly enhance efforts to prevent illicit use of the financial system and help reduce barriers to financial inclusion. While federated machine learning has been the object of much academic research,²⁷ there is a need for research to develop trustworthy systems for use in AFC risk management. Now is the time to engage with regulators, central banks, and global standard-setting bodies, to determine how to set policies for use of privacy-enhancing technologies like federated machine learning to enable firms do what the legal and regulatory infrastructure has long intended: learn from the experience of peer firms to make each individual firm's AFC processes and AFC detection across the system more effective.

We recommend that near-term research focus on understanding the relevant technologies and key policy issues, as the two topics are necessarily intertwined. This can be accomplished by:

- Evaluating existing federated learning approaches in AFC and adjacent applications like fraud with respect to their ability to improve how financial institutions distinguish those who can safely be given access to the financial system from those who cannot;
- Assessing trade-offs between performance, privacy, and inclusion for each form of federated learning in the research;
- Laying the foundations for development of interoperability standards to enable practicable sharing of model parameters; and
- Identifying policy options for central banks and other financial regulators to support and enable the adoption of technologies such as federated machine learning to improve the accuracy and efficiency of AFC processes (including KYC, CIP, and CDD as well as transaction monitoring).

²⁶ This kind of hub-and-spoke configuration for a federated learning system is not the only option, but it seems better suited on a preliminary basis to the AFC context than a peer-to-peer configuration, in which financial institutions share model updates directly with each other. The latter may offer less complete benefits than a hub-and-spoke configuration, since it would not enable firms to benefit from feedback on governmental investigation of SARs and other indications of illicit activity.

²⁷ Tian Li, Anit Kumar Sahu, Ameet Talwalkar, and Virginia Smith, Federated Learning: Challenges, Methods, and Future Directions, arXiv preprint arXiv:1908.07873, (Aug. 2019).



For evaluation of the technologies, we envision empirical analysis of several uses of federated learning technologies in AFC and adjacent applications like fraud detection. Such analysis would comprise experiments using large data sets (ideally real-world data containing known, and potentially unknown, instances of financial crimes) to train and test different federated learning systems. FinRegLab is in discussions with federated machine learning vendors and large banks about engaging in an evaluation of this type.

These federated learning systems would be assessed on the basis of their ability to detect illicit activity and would specifically focus on scenarios in which broad populations would likely be difficult to screen accurately under current methods. We would also evaluate the capacity of a federated learning system to protect sensitive information. Our research would also investigate potential issues and solutions regarding likely heterogeneity of models across different financial institutions that wish to collaborate on AFC risk detection, as well as the role a government entity could play as the hub in a federated learning system.

For the policy questions, we envision convening policy and regulatory stakeholders, as well as advocates, financial institutions, and technologists, to discuss the emerging technologies and the relevant regulations. These conversations would guide technology development and evaluation to ensure that both processes are being conducted in a manner that reflects policy purposes while informing the evolution of specific regulations to enable responsible innovation. In this process, it will also be important to consider the implications in emerging markets, where new approaches for establishing digital identity are being explored to help overcome infrastructure challenges to delivering public services and may point to better sources of information for AML customer screening that can be incorporated into a federated approach for screening customers and transactions for risk of illicit financial activity.

About FinRegLab

FinRegLab, Inc. is a not-for-profit research center that seeks to promote the safe use of new technology and data to improve the efficiency, stability, and fairness of the financial system. In light of technology-driven changes in financial services, public policy, and government have a core role to play to ensure that technology and data evolve and are used safely in driving financial inclusion. Policymakers and other stakeholders must have a practical and fact-based understanding of these new technologies and data, yet there is a dearth of evidence-based research and analysis to advance informed public dialogue and policy.

FinRegLab provides an independent, nonpartisan platform for policymakers, advocates, community organizations, and technology and financial sector firms to build an evidence-based understanding and assessment of new financial technologies and their impact on consumers and the implications for policy. We conduct empirical tests and policy research to inform and deepen the understanding of new financial technologies in a manner that is objective, fact-based, nonpartisan, fair, and balanced. FinRegLab also provides a platform for public engagement and dialogue to deepen stakeholders' understanding of the implications of technology and data in financial services. Our process features active engagement with key policymakers throughout the lifecycle of our projects in order to produce actionable insight that affects the evolution of policy.²⁸

²⁸ Consumer Compliance Outlook, Federal Reserve System (Third Issue 2019) (citing FinRegLab's research on the use of cash flow data in credit underwriting); Board of Governors of the Federal Reserve System, *et al.*, Interagency Statement on the Use of Alternative Data in Credit Underwriting (December 3, 2019).

