K2 Integrity

Promoting Financial Inclusion with Federated Learning:
The Consilient Solution

By Alex Levitov and Gail Fuller, K2 Integrity

Part I: The Challenge: Understanding and Managing the Risks of Underserved Populations

Whether it comes to pricing and extending credit, predicting fluctuations in the market, or protecting against operational breakdowns or abuse by illicit actors, financial institutions are in the business of managing risk. But in each of these domains, a bank's risk management decisions are only as good as the information and analytic tools they have at their disposal. Where such information is lacking—or, worse, where well-established ways of measuring risk are potentially biased or misleading—banks may shy away from doing business with certain customers or markets, leaving already underserved populations without vital access to financial services and missing out on business opportunities.

This phenomenon has manifested internationally over the past two decades, as many global banks have reassessed their risk exposure and tolerance regarding financial crime and sanctions. This has resulted in terminated or restricted business relationships with remittance companies and banks, primarily in developing or politically unstable countries perceived to be high risk for financial crime. Lacking a sufficiently granular and nuanced understanding of the risk that individual institutions and their customers present, global banks have at times opted to "de-risk" bluntly from entire regions, jurisdictions, or customer types, hindering global financial inclusion, reversing progress in reducing remittance prices and fees, and pushing higher-risk transactions into more opaque, informal channels.[1]

A similar dynamic is also undermining access to financial services closer to home. Relying on incomplete and often inaccurate tools for measuring creditworthiness and other risk factors, banks in the United States have long tended to avoid—rather than manage—the risks of providing services to the 53 million Americans who lack a credit score from a nationwide consumer reporting agency or the 56 million more whose credit scores are subprime.[2] Customers without a traditional credit score struggle to obtain mortgages, credit cards, or other bank loans and often must resort to high-cost loans that may not help build a credit history, even when successfully repaid.[3] Even where consumers have a traditional credit history, more than one in five has a "potentially material error" in their credit file.[4]

By the Numbers

53 MILLION

NUMBER OF AMERICANS WHO LACK A CREDIT SCORE FROM A NATIONWIDE CONSUMER REPORTING AGENCY

56 MILLION

ADDITIONAL AMERICANS WHOSE CREDIT SCORES ARE RATED SUBPRIME

1 IN 5

AMERICANS WHO HAVE A "POTENTIALLY MATERIAL ERROR" IN THEIR CREDIT FILE

The problem of inadequate credit scoring disproportionately impacts Black or Hispanic borrowers, who have not had equal access to homeownership and other sources of generational wealth, and those who live in low-income neighborhoods. It also impacts some recent immigrants, young people just getting started in the job market, and those who are recently widowed or divorced and lack a sufficient credit history on their own.[5] Moreover, even mission-driven banks such as community development financial institutions (CDFIs) and minority depository institutions (MDIs), which serve low-income and minority communities at higher rates than mainstream banks, may find it difficult to accurately risk-rate a largely "credit-invisible" customer base and price credit products accordingly—a challenge exacerbated as the COVID-19 pandemic has substantially weakened these banks' balance sheets and further strained profitability.[6]

In this paper, we outline how a new model K2 Integrity and Giant Oak pioneered with Consilient to help banks worldwide discover and manage their illicit finance risks could be deployed to enable small- and mid-sized banks in the United States—particularly CDFIs and MDIs—to better understand, assess, and manage the risks of serving customers who lack traditional credit scores or whose credit scores may inaccurately reflect their true credit risk. The **Consilient** solution is driven by an innovative technology and governance model known as "**federated learning**," whereby a machine-learning algorithm accesses and interrogates data sets across different institutions without ever moving or extracting the underlying data, thus enabling collective learning while safeguarding sensitive customer data.

> Federated learning solutions allow an algorithm to learn from diverse data sets without pooling or otherwise moving the data, facilitating collective learning without compromising the privacy of customer information.
>
> Because of their mission, CDFIs and MDIs likely have experiences successfully lending to "credit-invisible" populations or those that would be rated high risk using conventional means, presenting a critical opportunity for learning.

Supported by ongoing and expert human oversight and appropriate governance, this technology would learn from the behavior of customers across a consortium of CDFIs and MDIs to identify those who have payment histories that indicate they are low-risk from a credit perspective despite being "credit-invisible" or historically rated as higher risk. By learning from these cases, Consilient can help inform a more nuanced and accurate approach to risk scoring these customers and managing credit risk. Alternatives or complements to the traditional credit score-based model would allow CDFIs, MDIs, and other financial institutions to expand and more accurately price services to their existing customer base and reach new customers previously shut out of the formal financial system altogether.

Part II: The Consilient Solution: Secured Federated Learning

Consilient is the product of a partnership between K2 Integrity, the world's premier strategic advisory firm dedicated to financial integrity and security, and Giant Oak, a leading technology company that combines behavioral science and machine learning to build software solutions that combat illicit activity.[7] Originally developed to predict and detect money laundering, fraud, and other illicit finance threats to the financial system, Consilient's technology and federated governance framework can also be deployed to help governments and financial institutions manage a wider range of risks—uncovering not just bad actors attempting to abuse the financial system but legitimate, profitable customers inaccurately rated as "high risk."

At the core of Consilient's technology is the principle of federated machine learning, which allows different institutions to gain insights from one another's data, not by moving the data into a centralized repository, but by moving an analytic algorithm to the various local data environments. As the Consilient algorithm travels across a consortium of participating banks, it ingests, trains on, and learns from each institution's data—adjusting the model's weights and parameters as it is fitted to the local data, and sending encrypted updates to a secure aggregation engine without removing any data from its local source. As an added safeguard, both the data and the model are protected from theft, snooping, or tampering by Intel's Software Guard Extensions (Intel® SGX) technology, a hardware-based trusted execution environment that helps isolate and protect the confidentiality and integrity of code and data at each step of the federated learning process.[8]

Aggregating data insights via secured federated learning is particularly useful when trying to collaboratively analyze or verify sensitive financial data, such as data related to income verification or other markers of credit risk, as banks are naturally reluctant to reveal raw data to competitors and may be subject to data privacy or data localization rules that prevent them from transferring this information to other institutions or to a public cloud. The Consilient solution overcomes these barriers to information sharing by bringing the algorithm to the data, thus decoupling the ability to do machine learning from the need to store data in the cloud or in a single, centralized datacenter.

A Traveling Algorithm

Consilient helps overcome barriers to information sharing by bringing the algorithm to the data, thus decoupling the ability to do machine learning from the need to store data in the cloud or in a single, centralized data center.

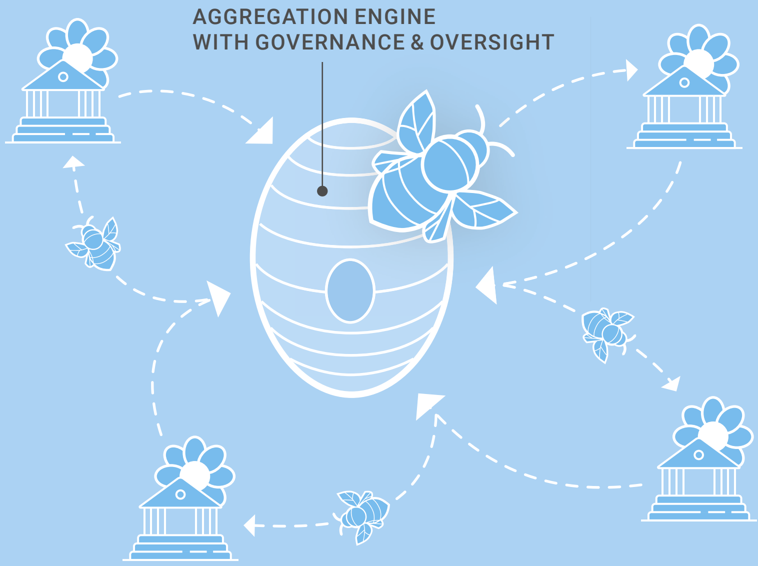

Part III: Advantages of Federated Learning for CDFIs and MDIs

Three features of the federated learning model make it especially well-suited to help CDFIs and MDIs safely and cost-effectively expand access to financial services.

- **The right partners:** Lacking insight into other institutions' data or their experience with comparable customer bases, CDFIs and MDIs are frequently left with an unattractive choice: either limit themselves to their own customers' past transaction histories as a basis for assessing future risks, or rely on a third-party credit scoring model that may be unreliable and even biased against the very populations these institutions are designed to serve. The federated learning alternative would allow CDFIs and MDIs from across the country to partner in a consortium of like-minded, mission-driven institutions, both expanding the pool of relevant customer data available for analysis and leveraging their unique understanding of the financial services in greatest demand—as well as the barriers to access historically faced—by low-income and minority communities.
- **The right training data:** Traditional credit scoring models rely on consumers' history of repaying debts such as mortgages and bank loans, together with records reflecting whether the consumer has any bills in collection or a history of liens, judgments, or bankruptcies. Yet these data sources paint at best an incomplete picture of a consumer's creditworthiness and neglect a range of other well-documented, measurable indicators of credit risk and past payment history. The Consilient algorithm would fill these critical data gaps by training on so-called "alternative data," including a consumer's history of paying monthly rent, utilities, and cell phone bills, as well as consumers' general management of deposit account cash balances, which can show a track record of meeting obligations that may not turn up in a traditional credit history.[9]

  Drawing on this expanded range of data sources—not just at a single institution but across the full consortium of participating banks—the machine-learning algorithm can identify risk attributes that may not have been weighted accurately in past credit decisions and apply the corrected model to the assessment of prospective customers and lending opportunities.[10] Just as importantly, the Consilient model is dynamic, meaning it updates continuously to account for new data and the evolving significance of individual data points, allowing banks to more accurately track consumers' changing behavior and financial circumstances. Traditional credit models, by contrast, often remain static until the model is periodically reviewed, refreshed with a new data set, and manually revised.[11]
- **A cost-sharing model:** Although there is widespread recognition among some experts that an improved risk-scoring solution presents a potentially transformative business opportunity for banks—allowing them to reach new customers and expand services through existing relationships—creating a new model from scratch requires time, data, technical expertise, and other resources that may not be available to most small- or mid-sized financial institutions. Indeed, resource constraints, as we have seen, may be especially acute at CDFIs and MDIs, which disproportionately serve low- and middle-income customers and have faced higher loan losses and lower profits than mainstream banks in the wake of the COVID-19 crisis.[12] Importantly, however, the federated learning model is also a cost-sharing model, allowing participating institutions not only to learn from one another's data but to pool expertise and resources in a way that dramatically reduces the cost of piloting such a solution alone.

Based on more and better data, informed by practitioner expertise, and driven by machine learning technology, the federated learning model would bring CDFIs and MDIs into an active partnership to improve their understanding of risks and expand financial access to creditworthy, profitable, yet traditionally underserved customer populations.

Part IV: Preventing Bias through Governance and Oversight

Governance and oversight are key elements of the Consilient model because we recognize that although the use of machine learning technology can help to overcome many of the gaps and inaccuracies in traditional credit scoring models—shortcomings that have disproportionately disadvantaged minority and low-income consumers—these same analytic tools, left to operate on their own, could also exacerbate existing bias and further entrench historical patterns of discrimination. For example, traditional metrics for assessing risk such as location of residence closely track racial and socio-economic categories due to historical patterns or policies of discrimination, such as housing segregation and mortgage “redlining.” In the absence of expert monitoring and intervention, machine learning algorithms may naturally assign risk-relevance to factors that, although apparently neutral, are in fact proxies for race or other “suspect variables.”[13] As research by the Brookings Institution’s experts observes, “Proxy discrimination by AI [artificial intelligence] is even more concerning because the machines are likely to uncover proxies that people had not previously considered.”[14]

It would clearly undermine the objectives and mission of CDFIs and MDIs if they were to utilize a machine learning algorithm that reproduced or enhanced—while perhaps also obscuring—the very biases and barriers to financial inclusion they are seeking to uproot. It is therefore essential that any application of the federated learning model be subject to rigorous governance and oversight mechanisms, including routine intervention by human experts to uncover proxy discrimination or other biases imposed or perpetuated by the technology.[15]

Following the governance framework Consilient has instituted for other federated learning consortia, appropriate subject-matter experts, both internal as well as those within the Consilient network, such as participating banks and any sponsoring government agency or non-governmental organization, would be consulted and engaged for governance and oversight, including: (i) reviewing and filtering data sets to prevent the consideration of prohibited attributes or known proxies for prohibited attributes; (ii) identifying and addressing clear indications that data sets are not representative of the relevant population; and (iii) regularly tuning the model to discourage the reproduction of bias. As such, the Consilient approach to governance would be informed by experts in credit risk and financial inclusion and supported by Consilient’s in-house technical experts to ensure ongoing oversight and the implementation of testing and assurance processes that would complement traditional fair lending testing and model testing and validation practices already under way at participating banks, ensuring that the new model can be explained and justified to supervisors.[16]

Part V: Conclusion

Implementing Consilient's federated learning model within a consortium of CDFIs and MDIs has the potential to transform the way financial institutions approach lending to "credit invisible" and historically underserved populations and to dramatically expand access to financial services.

By working with carefully selected data from CDFIs and MDIs who already serve these populations, implementing a federated model, and ensuring governance and oversight of the machine learning algorithm, the Consilient model helps to overcome many of the traditional barriers to implementing transformative new technologies in this space—a dearth of data on which to train machine learning models, data privacy, security, and competitive concerns, high costs, worries about unintentionally reinforcing existing biases or introducing new biases, and challenges associated with the explainability and defensibility of new approaches. This new approach has the potential to enable financial inclusion, security, and equity.

Our Team

Gail Fuller

Managing Director
gfuller@k2integrity.com ▸
Bio ▸

Alex Levitov

Senior Director
alevitov@k2integrity.com ▸
Bio ▸

1 The World Bank, "De-Risking in the Financial Sector," October 7, 2016, available at: https://www.worldbank.org/en/topic/financialsector/brief/de-risking-in-the-financial-sector.

2 AnnaMaria Andriotis, "Need Cash? Companies Are Considering Magazine Subscriptions and Phone Bills When Making Loans," The Wall Street Journal, September 12, 2019, available at: https://www.wsj.com/articles/need-cash-companies-are-considering-magazine-subscriptions-and-phone-bills-when-making-loans-11568280601?mod=article_inline.

3 Grovetta Gardineer, "Expanding Financial Inclusion for Credit Invisible Consumers and Businesses," Project REACh: A Look Back at the First Year, at pp. 5-6, available at: https://www.occ.treas.gov/topics/consumers-and-communities/minority-outreach/project-reach-anniversary.pdf.

4 Aaron Klein, "The Real Problem with Credit Reports Is the Astounding Number of Errors," September 28, 2017, available at: https://www.brookings.edu/research/the-real-problem-with-credit-reports-is-the-astounding-number-of-errors/.

5 Consumer Financial Protection Bureau, "CFPB Explores Impact of Alternative Data on Credit Access for Consumers Who Are Credit Invisible," February 16, 2017, available at: https://www.consumerfinance.gov/about-us/newsroom/cfpb-explores-impact-alternative-data-credit-access-consumers-who-are-credit-invisible/.

6 Federal Reserve Bank of San Francisco, "Minority Banks during the COVID-19 Pandemic," August 2, 2021, available at: https://www.frbsf.org/economic-research/publications/economic-letter/2021/august/minority-banks-during-covid-19-pandemic/.

7 K2 Integrity, "K2 Intelligence Financial Integrity Network and Giant Oak Collaborate with Intel to Launch Consilient," October 29, 2020, available at: https://www.k2integrity.com/en/newsroom/news-releases/k2-intelligence-fin-and-giant-oak-collaborate-with-intel-to-launch-consilient

8 For additional details, see Gary Shiffman, Juan Zarate, Nikhil Deshpande, Raghuram Yeluri, and Parviz Peiravi, "Federated Learning through Revolutionary Technology," available at: https://www.intel.com/content/www/us/en/financial-services-it/federated-learning-solution.html.

9 See Karan Kaul, "Adopting Alternative Data in Credit Scoring Would Allow Millions of Consumers to Access Credit," March 15, 2021, available at: https://www.urban.org/urban-wire/adopting-alternative-data-credit-scoring-would-allow-millions-consumers-access-credit. Preliminary findings from FinRegLab show that cash-flow underwriting outperforms traditional scoring models on its own, and when combined with traditional models is even more predictive; see FinRegLab, The Use of Cash-Flow Data in Underwriting Credit: Empirical Research Findings, July 2019, available at: https://finreglab.org/wp-content/uploads/2019/07/FRL_Research-Report_Final.pdf.

10 In addition to training on customer and transaction data from participating banks, the consortium could also partner with relevant government-sponsored entities, such as Fannie Mae or Freddie Mac, to obtain additional training data that could further improve the accuracy of the scoring model. On Fannie Mae and Freddie Mac's openness to alternative credit scoring models, see Andrew Ackerman, "Fannie, Freddie to Consider Alternatives to FICO Scores," The Wall Street Journal, August 13, 2019, available at: https://www.wsj.com/articles/fannie-freddie-to-consider-alternatives-to-fico-scores-under-new-rule-11565719353.

11 See Bank Policy Institute and Covington & Burling LLP, Artificial Intelligence: Recommendations for Principled Modernization of the Regulatory Framework, September 14, 2020, at 4, available at: https://bpi.com/wp-content/uploads/2020/10/Artificial-Intelligence-Recommendations-for-Principled-Modernization.pdf.

12 Federal Reserve Bank of San Francisco, "Minority Banks during the COVID-19 Pandemic," August 2, 2021, available at: https://www.frbsf.org/economic-research/publications/economic-letter/2021/august/minority-banks-during-covid-19-pandemic/.

12 Anya E.R. Prince and Daniel Schwarcz, "Proxy Discrimination in the Age of Artificial Intelligence and Big Data," Iowa Law Review, vol. 105 (2020), available at: https://ilr.law.uiowa.edu/assets/Uploads/ILR-105-3-Prince-Schwarcz-6.pdf.

12 Aaron Klein, "Reducing Bias in AI-Based Financial Services," July 10, 2020, available at: https://www.brookings.edu/research/reducing-bias-in-ai-based-financial-services/.

15 See Gary Shiffman, "We Need a New Field of AI to Combat Racial Bias," July 3, 2020, available at: https://techcrunch.com/2020/07/03/we-need-a-new-field-of-ai-to-combat-racial-bias/.

16 For additional examples of how AI credit underwriting systems could mirror standards that apply to conventional systems, see Bank Policy Institute and Covington & Burling LLP, Artificial Intelligence: Recommendations for Principled Modernization of the Regulatory Framework, September 14, 2020, at 17, available at: https://bpi.com/wp-content/uploads/2020/10/Artificial-Intelligence-Recommendations-for-Principled-Modernization.pdf.

“Do no harm, do the right thing, and do what you promised to do.”

Jules Kroll
Chairman and Co-Founder