



info@starlingtrust.com

September 22, 2020

via email: comments@fdic.gov

Robert E. Feldman,
Executive Secretary
Attention: Comments
Federal Deposit Insurance Corporation
550 17th Street, NW
Washington, DC 20429

RE: Request for Information on Standard Setting and Voluntary Certification for Models and Third-Party Providers of Technology and Other Services
Agency: FDIC
Docket ID: RIN 3064-ZA18

INTRODUCTION

Thank you for the opportunity to provide comments to the FDIC's Request for Information (RFI) related to Standard Setting and Voluntary Certification for Models and Third-Party Providers of Technology and Other Services. Starling (<https://starlingtrust.com>) is an innovative US-based RegTech startup that delivers analytics using internal bank data to improve non-financial risk governance, particularly with regard to risks that stem from firm culture.

Through our thought leadership and industry engagement, Starling has become recognized as an expert in our industry. Our annual white-paper, *Culture and Conduct Risk Management in the Banking Industry*¹, (aka the Starling 'Compendium'), has become a must-read reference on the latest trends and strategies taken by bank supervisors globally to address these non-financial operational risks.

Starling also offers an AI-driven technology platform that applies advances in behavioral science and network theory to the challenge of identifying and mitigating non-financial risk in banks – proactively.

We are strongly supportive of the FDIC's proposal to sponsor an industry certification program for emerging technology applications in the banking sector. As a leading startup in the RegTech space, we have faced numerous challenges in beginning work with banks interested in adopting our technology despite strong technical validation and use cases. A particular challenge has been the number of reviews that banks require in order to vet and validate technology before it is made available for the business to trial. We believe that a regulator-sponsored certification program could significantly reduce the cost and time involved for banks to explore the potential of new technologies and to experiment more frequently.

¹ <https://starlingtrust.com/compendium/>

BACKGROUND ON REGTECH SOLUTIONS FOR MANAGING OPERATIONAL RISK

The Basel Committee on Bank Supervision (“BCBS”) defines Operational Risk as the “risk of loss resulting from inadequate or failed internal processes, people and systems or from external events.”² For the past decade, spending on systems and processes to manage non-financial risk has exploded. Much of this was driven by legislative and regulatory changes implemented in the wake of the Financial Crisis and earlier scandals at firms like Enron. Banks have invested billions into processes and systems for governance, risk and compliance (GRC). Increasingly, intrusive surveillance and monitoring tools, often powered by AI, are gaining traction. At the same time, compliance and risk functions are increasingly turning to Robotic Process Automation (RPA) to replace manual tracking and reporting activities with automation.

At the operational level, banks have focused on implementing systems and processes to manage misconduct through controls, with a view to managing risk through documentation, restrictive processes, by removing people from decision making loops, and by detecting bad actors through surveillance and monitoring when controls fail – as they regularly do.

Without insight into the behavioral context that lies behind misconduct, management interventions are heavily rules-based, primarily targeting visible activities rather than underlying norms and cultural propensities. Further, by focusing on outcomes rather than the relational dynamics among teams that often precede *contagious* misconduct, standard non-financial risk management approaches are necessarily backward looking. Risk management becomes a tick-box exercise that is not ‘fit for purpose’ – amply evidenced by continual misconduct scandals.

Risk management functions that have oversight responsibilities, and manage to the bank’s risk appetite, have relied on management frameworks modeled on the 3 Lines of Defense (‘3LoD’) to manage the bank’s exposure to non-financial risk. Banks have implemented complex reporting systems and detailed processes to manage these frameworks. What these investments miss is the “people” piece of the puzzle called out by the BCBS. Yet the success of the 3LoD framework depends entirely upon a complex web of interactions and critical behaviors among senior executives and risk management specialists in order to function effectively. Unfortunately, tools like online surveys and townhall meetings do not adequately capture such complexity.

Many banking regulators rely on models based on the Uniform Financial Institution Rating System (“UFIRS”), also known as the ‘CAMELS ratings’, to generate an assessment of the overall bank that takes into account all of the significant financial and operational factors that represent a bank’s risk management practices. A fundamental element is the ‘Management Assessment’ (‘M’) component. Whereas other financial-related components can be supported by metrics and models, the M component remains largely subjective and imprecise. As a result, firms have difficulty knowing whether the systems and processes they have implemented are sufficient. It is also far more difficult to benchmark performance in these areas against peers through horizontal reviews. Furthermore, this opens regulators up to potential criticism when apparent inconsistencies between operations and assessments crop up.

² Sound Practices for the Management and Supervision of Operational Risk, BIS, February 2003

WHAT'S MISSING

In each of these cases, regulatory guidance strongly influences the decisions banks make in choosing the risk management framework they implement. Regulators and firms have prioritized processes and systems for internal risk governance (and guarding against external threats such as those in cybersecurity). They have been far less inclined to address the people element – namely how to foster the necessary behaviors and cultural norms required to manage those systems and processes correctly.

This is understandable because, for a long time, the tools available for measuring and managing behavior have not lent themselves to effective supervision or bank examination. Firms have been forced to rely on HR-delivered tools such as staff surveys, townhall meetings, self-reported behavior journals, and online ethics training. These tools lack objectivity, specificity, and real-time responsiveness. And – when these measures fail – the fallback is reliance upon robust surveillance and monitoring systems that promise to detect risk events as they occur. Such instruments produce high numbers of ‘false-positive’ signals which result in added expense as risk examiners are required to run each to ground. And, when successful in identifying an actual risk management failure, awareness of such is too little / too late.

These challenges are all the more relevant in the current COVID-19 pandemic. Controls and surveillance systems that were established in a time when everyone worked together have been upended. Further, the most effective protection is provided by a culture that encourages challenge and speak-up behavior, and where staff feels able and encouraged to push back the moment that risk behaviors threaten to take hold.

In a work-from-home environment, this too has been severely weakened. During a recent interview with *Bloomberg*, Gary Cohn, past-COO of Goldman Sachs and advisor to Starling, was quoted as saying: “Banks need people to be working together in a cooperative fashion and watching and listening to each other,” adding, “That is what the Fed would call a first line of defense: overhearing conversations, looking at presentations, or looking at the way you talk to a client. [...] When people are sitting in their bedrooms, there is no one there to look over their shoulder.”³

This situation will not be solved by existing approaches. Rather, banks need to test new technologies, models, and frameworks that can serve to break this impasse. Regulators play a key role in this, as ongoing bank scandals contribute to an erosion in the public’s faith in ‘the system.’ By promoting innovation in risk management, regulators work to protect/promote the public’s interests.

MACHINE LEARNING OFFERS A WAY FORWARD

Advances in machine learning have made it possible to sift through vast troves of internal bank data at scale. By applying novel approaches in the field of “computational social science,” it is now possible to detect signals within those massive data sets that tie to particular behaviors of interest to management and supervisors. These may be behaviors that represent a predilection for misconduct or, equally, behaviors that are necessary to the full functioning of critical non-financial risk management systems and processes.

Analyzing these signals allows us to generate metrics that update continuously and reveal where specific behavioral propensities are likely to appear. Such tools can illuminate the pathways by which certain behaviors are most likely to spread – contagion-like – throughout an organization. This ‘behavioral

³ <https://www.bloomberg.com/opinion/articles/2020-07-08/covid-19-pandemic-is-a-great-incubator-for-financial-fraud?sref=GNTXiFne>

epidemiology' positions management to operate from the front-foot. It also allows precision targeting of audit activities and risk management interventions, allowing firms and supervisors to scale their risk oversight and to act in a more timely, effective, and efficient manner.

A significant additional benefit is to be had once such technologies are established as industry-standard best practice: standardized risk metrics such as those we describe here may permit for horizontal reviews on an apples-to-apples basis, system-wide, across any given jurisdictional space. And the adoption of such metrics among regulators in other financial markets may permit for more efficient collaborative oversight of firms across their global footprint.

RECOMMENDATIONS

Re: Questions regarding potential advantages to service providers and IDIs

There are a number of challenges to the adoption of innovative technologies that could be addressed through a standard-setting and voluntary certification process.

In our engagement with banks across the globe, we have experienced a consistent lack of expertise in the ability to evaluate potential *regulatory* risks associated with the adoption of new technologies and models. The pace of change and sophistication required to maintain expertise across disciplines is simply too great. This holds for large global institutions as well as for small, regional banks. To the extent that a bank even has a formal process, a potential vendor must deal with multiple, often redundant, rounds of reviews by various stakeholders. We have seen several cases where this process is so challenging that even large banks have delayed consideration of promising technologies because of resource constraints.

Trialing new tools under a certification system would bring needed structure to initiatives that require collective action across the industry. Firms struggle to achieve such collective action in the absence of more formal industry platforms that can provide a forum for such action. A standard-setting process and certification program that brought together multiple industry stakeholders would facilitate collaborative engagement between regulators, firms, and technology vendors. This alone would help to engender a more meaningful dialogue across the ecosystem which would be a benefit in and of itself.

By providing support to a certification process, regulators like the FDIC can encourage adoption of new technologies. Particularly in the case of operational risk where innovation has lagged, banks may lack confidence as to how new technologies may offer value. As a result, banks are incentivized to simply do more of the same. A certification process could provide an avenue by which a regulator could signal that emerging technologies could be useful while still falling well short of an explicit endorsement of such technology. This too would encourage banks to try new technologies that go beyond well-known, but marginally effective, solutions.

Re: Questions regarding Scope

Standard setting and certification should avoid being prescriptive around models, algorithms, and related services. Many promising technologies are still in a relatively early stage of development and it is important to avoid standards that might inadvertently restrict desired innovation. Instead, the FDIC can sponsor standards that focus on setting and promoting industry best practices. This can be further

complemented through increasing disclosure requirements to help banks make responsible decisions about the technologies they are evaluating.

A full analysis of these technologies should be done by an appropriate and qualified stakeholder group. We see three primary areas where a certification program could focus: collecting and managing sensitive data, responsible algorithms, and appropriate use of model outcomes.

Collecting and Managing Sensitive Data

Technology solutions, particularly those reliant on machine-learning, often rely on sensitive data sets such as personally identifiable information (PII), email or message content, and location data. Banks have an obligation to protect such data. This concern alone drives a significant degree of scrutiny related to data privacy, security, and appropriate use.

We would recommend that future standards take the approach that sensitive data (e.g. email content, transaction data) should be used only to the extent it is necessary. Furthermore, there should be an expectation of disclosure when such data is used so that those subjected to such data collection, namely customers and employees, can be made aware how this data is being used.

Responsible Algorithms

Machine Learning algorithms can be very complex which can make it difficult to interpret the resulting outcomes or recommendations. As a result, there are a number of ways in which biases can be introduced into the data set, often in subtle ways. These biases can insert themselves through the choice of data sets or the structure of the model itself.

For example, employees of a bank may be monitored to detect patterns of activity associated with good management. However, if the model of ‘good management’ is based on past managers that have been predominately male, then the algorithm may be more likely to identify traits associated with men going forward. This kind of bias has been discovered in many applications including criminal detection and credit verification.

Even large banks may lack the technical capability and bandwidth to properly evaluate external algorithms. At the same time, it would be difficult to establish specific standards around model development in a way that would still encourage innovation. Establishing industry best practices for managing bias can be coupled with disclosure requirements as to how such practices are applied can address this risk. Such guidance can also help to educate bank employees as to how these algorithms and models work along with their limitations.

Appropriate Use of Model Outcomes

Models can be very powerful, but they come with limitations. It is rarely effective or appropriate for banks to directly execute on recommendations generated by these models. Far better outcomes can be achieved by pairing model recommendations with human judgement. A standard-setting and certification process should be designed with this approach in mind.