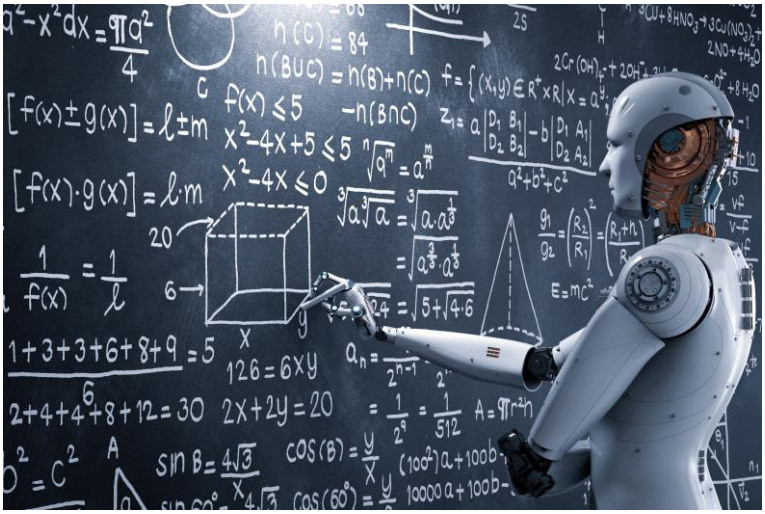

ARTIFICIAL INTELLIGENCE & MACHINE LEARNING

A Primer

ABSTRACT

To save you the trouble and embarrassment of raising your hand when asked at a conference “is there anyone out there that needs me to explain what artificial intelligence and machine learning are?”, here’s a short primer made up of three articles written by subject matter experts on ML and AI, and six articles written by a subject matter enthusiast …

James Richards

ACFCS Seminar – December 5, 2019

Table of Contents

Artificial Intelligence, Machine Learning, Deep Learning – A Primer

As we get closer to the year 2020, no one wants to admit that they don't really understand what artificial intelligence is. Or the difference between artificial intelligence and machine learning. After all, these terms have been around for years, and the "experts" lob them about as if everyone knows what they mean.

To save you the trouble and embarrassment of raising your hand when asked at a conference "is there anyone out there that needs me to explain what artificial intelligence and machine learning are?", here's a short primer.

The first explanations come from Meenal Dhande from an article that dated May 6, 2017 from https://www.geospatialworld.net/blogs/difference-between-ai%EF%BB%BF-machine-learning-and-deep-learning/

The Difference Between Machine Learning and Deep Learning - Meenal Dhande

You can think of artificial intelligence (AI), machine learning and deep learning as a set of a *matryoshka* doll, also known as a Russian nesting doll. Deep learning is a subset of machine learning, which is a subset of AI.

Artificial intelligence is any computer program that does something smart. It can be a stack of a complex statistical model or if-then statements. AI can refer to anything from a computer program playing chess, to a voice-recognition system like Alexa. However, the technology can be broadly categorized into three groups — Narrow AI, artificial general intelligence (AGI), and superintelligent AI.

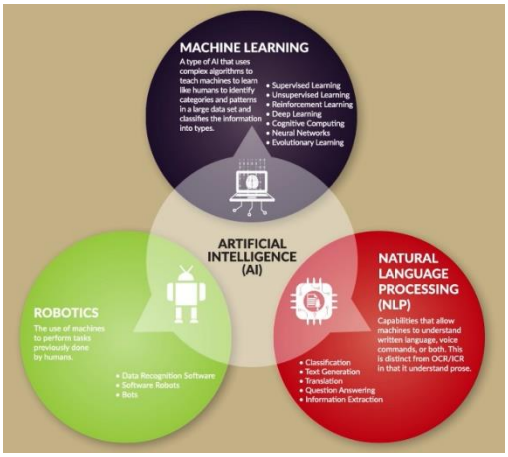

IBM's Deep Blue, which beat chess grandmaster Garry Kasparov at the game in 1996, or Google DeepMind's AlphaGo, which beat Lee Sedol at Go in 2016, are examples of narrow AI — AI that is skilled at one specific task. This is different from AGI — AGI is the intelligence of a machine that could successfully perform a range of tasks intellectual task that a human being can. On the other hand, Superintelligent AI takes things a step further. As Nick Bostrom describes it, this is "an intellect that is much smarter than the best human brains in practically every field, including scientific creativity, general wisdom, and social skills." In other words, it is when the machines have outfoxed us.

Machine learning is a subset of AI. The theory is simple, machines take data and 'learn' for themselves. It is currently the most promising tool in the AI pool for businesses. Machine learning systems can quickly apply knowledge and training from large datasets to excel at facial recognition, speech recognition, object recognition, translation, and many other tasks. Machine learning allows a system to learn to recognize patterns on its own and make predictions, contrary to hand-coding a software program with specific instructions to complete a task. While Deep Blue and DeepMind are both types of AI, Deep Blue was rule-based, dependent on programming — so it was not a form of machine learning. DeepMind, on the other

hand — beat the world champion in Go by training itself on a large data set of expert moves. That is, all machine learning counts as AI, but not all AI counts as machine learning.

Deep learning is a subset of machine learning. Deep artificial neural networks are a set of algorithms reaching new levels of accuracy for many important problems, such as image recognition, sound recognition, recommender systems, etc. It uses some machine learning techniques to solve real-world problems by tapping into neural networks that simulate human decision-making. Deep learning can be costly and requires huge datasets to train itself. This is because there are a huge number of parameters that need to be understood by a learning algorithm, which can primarily yield a lot of false-positives. For example, a deep learning algorithm could be trained to 'learn' how a dog looks like. It would take an enormous dataset of images for it to understand the minor details that distinguish a dog from a wolf or a fox.

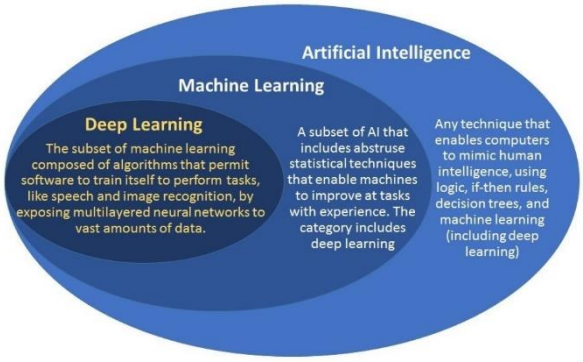

The second comes from Karen Hao, who (brilliantly and humorously) writes about technology for the MIT Technology Review. On November 17, 2018, Ms. Hao answered the question "what is the definition of machine learning?" at https://www.technologyreview.com/s/612437/what-is-machine-learning-we-drew-you-another-flowchart/

What is the definition of machine learning? – Karen Hao

Machine-learning algorithms use statistics to find patterns in massive amounts of data. And data, here, encompasses a lot of things—numbers, words, images, clicks, what have you. If it can be digitally stored, it can be fed into a machine-learning algorithm. Machine learning is the process that powers many of the services we use today—recommendation systems like those on Netflix, YouTube, and Spotify; search engines like Google and Baidu; social-media feeds like Facebook and Twitter; voice assistants like Siri and Alexa. The list goes on.

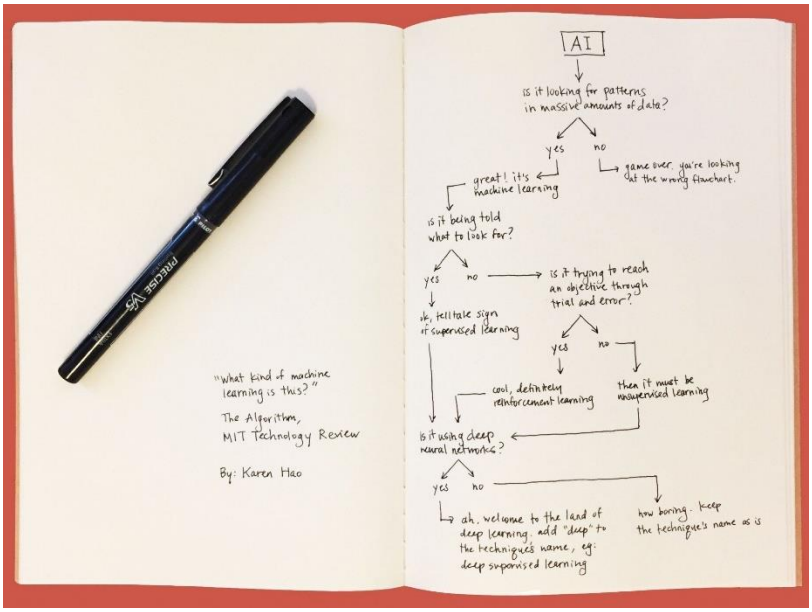

In all of these instances, each platform is collecting as much data about you as possible—what genres you like watching, what links you are clicking, which statuses you are reacting to—and using machine learning to make a highly educated guess about what you might want next. Or, in the case of a voice assistant, about which words match best with the funny sounds coming out of your

mouth. Frankly, this process is quite basic: find the pattern, apply the pattern. But it pretty much runs the world. That's in big part thanks to an invention in 1986, courtesy of Geoffrey Hinton, today known as the father of deep learning.

What is deep learning?

Deep learning is machine learning on steroids: it uses a technique that gives machines an enhanced ability to find—and amplify—even the smallest patterns. This technique is called a deep neural network—deep because it has many, many layers of simple computational nodes that work together to munch through data and deliver a final result in the form of the prediction.

What are Neural Networks?

Neural networks were vaguely inspired by the inner workings of the human brain. The nodes are sort of like neurons, and the network is sort of like the brain itself. (For the researchers among you who are cringing at this comparison: Stop pooh-poohing the analogy. It's a good analogy.) But Hinton published his breakthrough paper at a time when neural nets had fallen out of fashion. No one really knew how to train them, so they weren't producing good results. It took nearly 30 years for the technique to make a comeback. And boy, did it make a comeback.

What is supervised learning?

One last thing you need to know: machine (and deep) learning comes in three flavors: supervised, unsupervised, and reinforcement. In supervised learning, the most prevalent, the data is labeled to tell the machine exactly what patterns it should look for. Think of it as something like a sniffer dog that will hunt down targets once it knows the scent it's after. That's what you're doing when you press play on a Netflix show—you're telling the algorithm to find similar shows.

What is unsupervised learning?

In unsupervised learning, the data has no labels. The machine just looks for whatever patterns it can find. This is like letting a dog smell tons of different objects and sorting them into groups with similar smells. Unsupervised techniques aren't as popular because they have less obvious applications. Interestingly, they have gained traction in cybersecurity.

What is reinforcement learning?

Lastly, we have reinforcement learning, the latest frontier of machine learning. A reinforcement algorithm learns by trial and error to achieve a clear objective. It tries out lots of different things and is rewarded or penalized depending on whether its behaviors help or hinder it from reaching its objective. This is like giving and withholding treats when teaching a dog a new trick. Reinforcement learning is the basis of Google's AlphaGo, the program that famously beat the best human players in the complex game of Go.

The third explanation, or answer to the question "what is AI, exactly?" also comes from Karen Hao. This one was posted on November 10, 2018 at https://www.technologyreview.com/s/612404/is-this-ai-we-drew-you-a-flowchart-to-work-it-out/

What is AI, exactly? - Karen Hao

The question may seem basic, but the answer is kind of complicated.

In the broadest sense, AI refers to machines that can learn, reason, and act for themselves. They can make their own decisions when faced with new situations, in the same way that humans and animals can.

As it currently stands, the vast majority of the AI advancements and applications you hear about refer to a category of algorithms known as machine learning (see "What is machine learning?"). These algorithms use statistics to find patterns in massive amounts of data. They then use those patterns to make predictions on things like what shows you might like on Netflix, what you're saying when you speak to Alexa, or whether you have cancer based on your MRI.

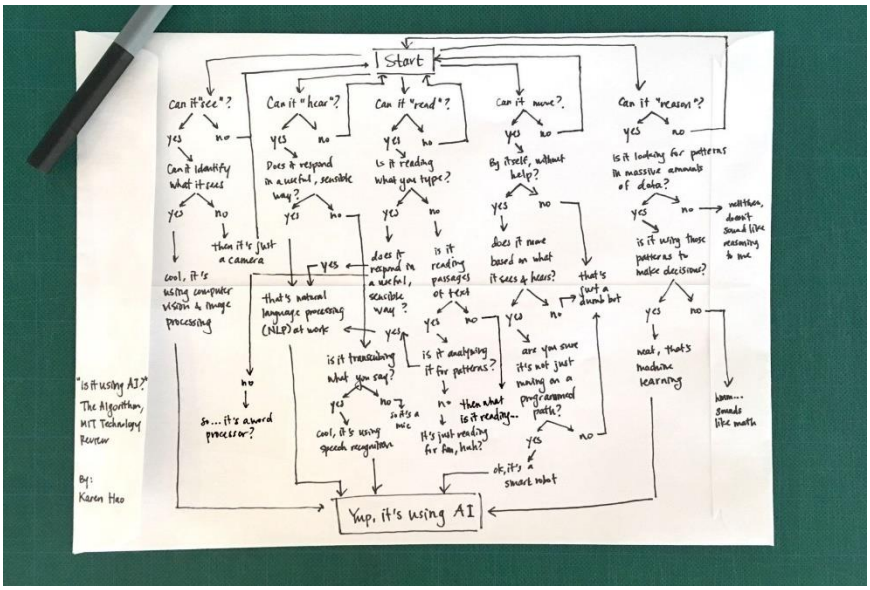

Machine learning, and its subset deep learning (basically machine learning on steroids), is incredibly powerful. It is the basis of many major breakthroughs, including facial recognition, hyper-realistic photo and voice synthesis, and AlphaGo, the program that beat the best human player in the complex game of Go. But it is also just a tiny fraction of what AI could be.

The grand idea is to develop something resembling human intelligence, which is often referred to as "artificial general intelligence," or "AGI." Some experts believe that machine learning and deep learning will eventually get us to AGI with enough data, but most would agree there are big missing pieces and it's still a long way off. AI may have mastered Go, but in other ways it is still much dumber than a toddler.

In that sense, AI is also aspirational, and its definition is constantly evolving. What would have been considered AI in the past may not be considered AI today.

Because of this, the boundaries of AI can get really confusing, and the term often gets mangled to include any kind of algorithm or computer program. We can thank Silicon Valley for constantly inflating the capabilities of AI for its own convenience. (Cough, Mark Zuckerberg, cough.)

Richards Articles on FinTech and Financial Crimes Risk Management

The following are articles written by James Richards and posted on www.regtechconsulting.net/news that deal with fintech innovation more broadly as it is applied for financial crimes risk management.

A Bank's Bid for Innovative AML Solutions: Innovation Remains A Perilous Endeavor

One Bank Asked the OCC to Have an "Agile Approach to Supervisory Oversight"

Posted November 11, 2019

On September 27, 2019 the OCC published an Interpretive Letter answering an unknown bank's request to make some innovative changes

to how it files cash structuring SARs. Tacked onto its three technical questions was a request by the bank to do this innovation along with the OCC itself through something the bank called an "agile approach to supervisory oversight." After qualified "yes" answers to the three technical questions, the OCC's Senior Deputy Comptroller and Chief Counsel indicated that the OCC was open to "an agile and transparent supervisory approach while the Bank is building this automated solution" but he didn't actually write that the OCC would, in fact, adopt an agile approach. This decision provides some insight, and perhaps the first public test, of (i) the regulators' December 2018 statement on using innovative efforts to fight money laundering, and (ii) the OCC's April 2019 proposal around innovation pilot programs. Whether the OCC passed the test is open to discussion: what appears settled, though, is that AML innovation in the regulated financial sector remains a perilous endeavor.

Regulators' December 2018 Joint Statement on Innovative AML Efforts

On December 3, 2018 the five main US Bank Secrecy Act (BSA) regulators issued a joint statement titled "Innovative Efforts to Combat Money Laundering and Terrorist Financing".[1] The intent of the statement was to encourage banks to use modern-era technologies to bolster their BSA/AML compliance programs. The agencies asked banks "to consider, evaluate, and, where appropriate, responsibly implement innovative approaches to meet their Bank Secrecy Act/anti-money laundering (BSA/AML) compliance obligations, in order to further strengthen the financial system against illicit financial activity" and "[t]he Agencies recognize[d] that private sector innovation, including new ways of using existing tools or adopting new technologies, can help banks" to do so.

The statement was a very positive step to encourage private sector innovation in fighting financial crime by testing new ways of using existing tools as well as adopting new technologies.

But it wasn't the "green light to innovate" that some people have said it is. There was some language in the statement that made it, at best, a cautionary yellow light. And the September 27th OCC letter seems to clarify that banks can innovate, but the usual regulatory oversight and potential sanctions still apply.

The Agencies' December 2018 statement included five things that bear repeating:

1. "The Agencies recognize that private sector innovation, including new ways of using existing tools or adopting new technologies, can help banks identify and report money laundering, terrorist financing, and other illicit financial activity by enhancing the effectiveness and efficiency of banks' BSA/AML compliance programs. To assist banks in this effort, *the Agencies are committed* to continued engagement with the private sector and other interested parties."
2. *"The* Agencies *will not penalize or criticize* banks that maintain effective BSA/AML compliance programs commensurate with their risk profiles but choose not to pursue innovative approaches."
3. "While banks are expected to maintain effective BSA/AML compliance programs, *the* Agencies *will not advocate* a particular method or technology for banks to comply with BSA/AML requirements."
4. Where test or implemented "artificial intelligence-based transaction monitoring systems … identify suspicious activity that would not otherwise have been identified under existing processes, the Agencies *will assess* the adequacy of banks' existing suspicious activity monitoring processes independent of the results of the pilot program"

5. "… the implementation of innovative approaches in banks' BSA/AML compliance programs *will not* result in additional regulatory expectations."

Note the strong, unqualified language: "the Agencies *are committed* to continued engagement", "the Agencies *will not* penalize or criticize", "the Agencies *will not* advocate …", "the Agencies *will* assess", and "the implementation of innovative approaches *will not* result in additional regulatory expectations".

The qualified "assurances" come in the paragraph about pilot programs (with emphasis added):

"Pilot programs undertaken by banks, in conjunction with existing BSA/AML processes, are an important means of testing and validating the effectiveness of innovative approaches. While the Agencies may provide feedback, pilot programs in and of themselves *should not* subject banks to supervisory criticism even if the pilot programs ultimately prove unsuccessful. Likewise, pilot programs that expose gaps in a BSA/AML compliance program *will not necessarily result in supervisory action* with respect to that program. For example, when banks test or implement artificial intelligence-based transaction monitoring systems and identify suspicious activity that would not otherwise have been identified under existing processes, the Agencies *will not automatically assume* that the banks' existing processes are deficient. In these instances, the Agencies will assess the adequacy of banks' existing suspicious activity monitoring processes independent of the results of the pilot program. Further, the implementation of innovative approaches in banks' BSA/AML compliance programs will not result in additional regulatory expectations."

Here there are the qualified assurances (a qualified assurance is not an assurance, by the way): "should not" is different than "will not"; "will not necessarily" is very different than "will not"; and "not automatically assume" isn't the same as "not assume". These are important distinctions. The agencies could have written something very different:

"… pilot programs in and of themselves *will not* subject banks to supervisory criticism even if the pilot programs ultimately prove unsuccessful. Likewise, pilot programs that expose gaps in a BSA/AML compliance program *will not result in supervisory action* with respect to that program. For example, when banks test or implement artificial intelligence-based transaction monitoring systems and identify suspicious activity that would not otherwise have been identified under existing processes, the *Agencies will not assume* that the banks' existing processes are deficient …"

The OCC's April 2019 Innovation Pilot Program

On April 30, 2019 the OCC sought public comment on its proposed Innovation Pilot Program, a voluntary program designed to provide fintech providers and financial institutions "with regulatory input early in the testing of innovative activities that could present significant opportunities or benefits to consumers, businesses, financial institutions, and communities." See OCC Innovation Pilot Program. As the OCC has written, the Innovation Pilot Program clearly notes that the agency would not provide "statutory or regulatory waivers and does not absolve entities participating in the program from complying with applicable laws and regulations."

Twenty comments were posted to the OCC's website. A number of them included comments that innovators needed some formalized regulatory forbearance in order to be able encourage them to

innovate. The Bank Policy Institute's letter (BPI Comment), submitted by Greg Baer (a long-standing and articulate proponent of reasonable and responsible regulation), provided that:

"… the OCC should clarify publicly that a bank is not required to seek the review and approval of its examination team prior to developing or implementing a new product, process, or service; that unsuccessful pilots will not warrant an MRA or other sanction unless they constitute and unsafe and unsound practice or a violation of law; and that innovations undertaken without seeking prior OCC approval will not be subject to stricter scrutiny or a 'strict liability' regime. We also recommend that the OCC revisit and clarify all existing guidance on innovation to reduce the current uncertainty regarding the development of products, processes and services; outdated or unnecessary supervisory expectations should be rescinded."

The American Bankers Association comment ABA Comment also asks for similar guidance:

"For institutions to participate confidently in a pilot, there must be internal agreement that OCC supervision and enforcement will not pursue punitive actions. In other words, the program should produce decisions that have the full support of the OCC and bind the agency to those conclusions going forward … One way for the OCC to accomplish this is to clarify that a participating bank will not be assigned Matters Requiring Attention (MRAs) if it acts in good faith as part of a Pilot Program. The nature of technological innovation means that banks must try new things, experiment, and sometimes make mistakes. The Pilot Program has been designed as a short-term limited-scale test to ensure that any mistakes made are unlikely to have an impact on the safety and soundness of an institution. Clarifying that MRAs will not be issued for mistakes made in good faith may help give banks the certainty they need to participate in a Pilot Program."

And the Securities Industry and Financial Markets Association (SIFMA) comment letter SIFMA Comment Letter included the following:

"Relief from strict regulatory compliance is a vital prerequisite to draw firms into the test environment, precisely so that those areas of noncompliance may be identified and remediated and avoid harm to the consumers. Without offering this regulatory relief, the regulatory uncertainty associated with participating in the Pilot Program could, by itself, deter banks from participating. Similarly, the lack of meaningful regulatory relief could limit the opportunity the program provides for firms to experiment and innovate."

So where did that leave banks that were thinking of innovative approaches to AML? For those that choose not to pursue innovative pilot programs, it is clear that they will not be penalized or criticized, but for those that try innovative pilot programs that ultimately expose gaps in their BSA/AML compliance program, the agencies will not *automatically* assume that the banks' existing processes are deficient. In response to this choice – do not innovate and not be penalized, or innovate and risk being penalized – many banks have chosen the former. As a result, advocates for those banks – the BPI and ABA, for example – have asked the OCC to clarify that it will not pursue punitive actions against banks that unsuccessfully innovate.

How has the OCC replied? It hasn't yet finalized its Innovation Program, but it has responded to a bank's request for guidance on some innovative approaches to monitoring for, alerting on, and filing suspicious activity reports on activity and customers that are structuring cash transactions.

A Bank's Request to Have the OCC Help It Innovate

The OCC published an Interpretive Letter on September 27, 2019 that sheds some light on how it looks at its commitments under the December 2018 innovation statement.[2] According to the Interpretive Letter, on February 22, 2019 an OCC-regulated bank submitted a request to streamline SARs for potential structuring activity (the Bank also sought the same or a similar ruling from FinCEN: as of this writing, FinCEN has not published a ruling). The bank asked three questions (and the OCC responded):

1. Whether the Bank could file a structuring SAR based solely on an alert, without performing a manual investigation, and if so, under what circumstances (yes, but with some significant limitations);
2. Whether the proposed automated generation of SAR narratives for structuring SARs was consistent with the OCC's SAR regulations (yes, but with some significant limitations);
3. Whether the proposed automation of SAR filings was consistent with the OCC's BSA program regulations (yes, but with some significant limitations).

The most interesting request by the Bank, though, was its request that the OCC take an "agile approach to supervisory oversight" for the bank's "regulatory sandbox" initiative. Pages 6 and 7 of the OCC letter provide the particulars of this request. There, the OCC writes:

"Your letter also requested regulatory relief to conduct this initiative within a "regulatory sandbox." Your regulatory sandbox request states 'This relief would be in the form of an agile approach to supervisory oversight, which would include the OCC's full access, evaluation, and participation in the initiative development, but would not include regulatory outcomes such as matters requiring attention, violations of law or financial penalties. [The Bank] welcomes the OCC to consider ways to participate in reviewing the initiative outcomes outside of its standard examination processes to ensure effectiveness and provide feedback about the initiative development.'"

NOTE: I had to read the key sentence a few times to settle on its intent and meaning. That sentence is "This relief would be in the form of an agile approach to supervisory oversight, which would include the OCC's full access, evaluation, and participation in the initiative development, but would not include regulatory outcomes such as matters requiring attention, violations of law or financial penalties."

Was the bank saying the relief sought was an agile approach to supervisory oversight that included the OCC's full participation in the process and no adverse regulatory outcomes? Or was the bank saying the relief sought was an agile approach to supervisory oversight that included the OCC's full participation in the process, but did not include anything to do with adverse regulatory outcomes?

I settled on the latter meaning: that the bank was seeking the OCC's full participation, but did not expect any regulatory forbearance.

The OCC first reiterated its position from the December 2018 joint statement by writing that it "supports responsible innovation in the national banking system that enhances the safety and soundness of the federal banking system, including responsibly implemented innovative approaches to meeting the compliance obligations under the Bank Secrecy Act." It then wrote that it "is also open to an agile and transparent supervisory approach while the Bank is building this automated solution for filing Structuring SARs and conducting user acceptance testing." This language is a bit different than what the OCC wrote at

the top of page 2 of the letter: "the OCC is open to engaging in regular discussions between the Bank and appropriate OCC personnel, including providing proactive and
timely feedback relating to this automation proposal."

Notably, the OCC wrote that it is "open to an agile and transparent supervisory approach", and "open to engaging in regular discussions between the Bank and appropriate OCC personnel", but being open to something doesn't mean you approve of it or agree to it. In fact, the OCC didn't appear to grant the bank's request. In the penultimate sentence the OCC wrote: "The OCC will monitor any such changes through its ordinary supervisory processes."

How About Forbearance to Innovate Without Fear of Regulatory Sanctions?

As set out above, in June 2019 the BPI and ABA (and eighteen others) commented on the OCC's proposal for an innovation pilot program. The BPI commented that "the OCC should clarify publicly that … unsuccessful pilots will not warrant an MRA or other sanction unless they constitute and unsafe and unsound practice or a violation of law", and the ABA commented that the OCC should "clarify that a participating bank will not be assigned Matters Requiring Attention (MRAs) if it acts in good faith as part of a Pilot Program".

The OCC seems to have obliquely responded to both of those comments. In its September 2019 Interpretative Letter, the OCC took the time to write that it "will not approve a regulatory sandbox that includes forbearance on regulatory issues for the Bank's initiative for the automation of Structuring SAR filings." Note that the OCC made this statement even though the bank appears to have specifically indicated that the requested relief did ***not*** include forbearance from "regulatory outcomes such as matters requiring attention, violations of law or financial penalties". And the OCC letter includes a reference to both the Interagency statement on responsible innovation and the OCC's April 2019 Innovation Pilot Program (see footnote 25 on page 7): "banks must continue to meet their BSA/AML compliance obligations, as well as ensure the ongoing safety and soundness of the bank, when developing pilot programs and other innovative approaches."

So although the OCC hasn't formally responded to the comments to its June 2019 innovation program to allow banks to innovate without fear of regulatory sanction if that innovation doesn't go well, it has made it clearer that a bank still has the choice to not innovate and not be penalized, or to innovate and risk being penalized.

(In fairness, in its Spring 2019 Semiannual Risk Perspective Report, the OCC noted that a bank's inability to innovate is "a source of significant strategic risk." See OCC Semiannual Risk Perspective, 2019-49 (May 20, 2019)).

Timely Feedback – Is Seven Months Timely?

As set out above, the OCC wrote that it "is open to engaging in regular discussions between the Bank and appropriate OCC personnel, including providing proactive and timely feedback …". The bank's request was submitted on February 22, 2019. The OCC's feedback was sent on September 27, 2019. So it took the OCC seven months to respond to the bank's request for an interpretive letter. In this age of high-speed fintech disruption, seven months should not be considered "timely." What would be timely? I would aim for 90 days.

Conclusion

This unnamed OCC-regulated bank appears to have a flashing green or cautionary yellow light from the OCC to deploy some technology and process enhancements to streamline a small percentage if its SAR monitoring, alerting, and filing. The OCC will remain vigilant, however, warning the bank that it "must ensure that it has developed and deployed appropriate risk governance to enable the bank to identify, measure, monitor, and control for the risks associated with the automated process. The bank also has a continuing obligation to employ appropriate oversight of the automated process."

So the message to the 1,700 or so OCC banks appears to be this: there's no peril in not innovating, but if you decide to innovate, do so at your peril.

[1] The Board of Governors of the Federal Reserve System, the Federal Deposit Insurance Corporation (FDIC), the Financial Crimes Enforcement Network (FinCEN), the National Credit Union Administration, and the Office of the Comptroller of the Currency. The statement is available at https://www.occ.gov/news-issuances/news-releases/2018/nr-occ-2018-130a.pdf

[2] https://www.occ.gov/topics/charters-and-licensing/interpretations-and-actions/2019/int1166.pdf

BigTech, FinTech, and the Battle Over Financial Services

BigTech vs FinTech – Which Will Replace Traditional Banks?

Posted May 26, 2019

Two recent papers have looked at the attributes, relative strengths and weaknesses, and likelihood to emerge as the main challenger to traditional financial institutions, of two different species of technology company: BigTechs and FinTechs. The two papers are:

- Financial Stability Board's (FSB) February 2019 paper titled "FinTech and Market Structure in Financial Services", available at https://www.fsb.org/wp-content/uploads/P140219.pdf
- Bank for International Settlements' (BIS) April 2019 Working Paper titled "BigTech and the changing structure of financial intermediation", available at https://www.bis.org/publ/work779.pdf

The BIS Working Paper makes a pretty compelling argument that the BigTech firms have some distinct advantages over FinTechs that make them more likely to usurp traditional financial institutions. Advantages such as an existing customer base (that is familiar with a user interface and messaging platform), and access to capital (often without the constraints that financial institutions have). And the BIS paper also sets out some of the advantages that BigTech has over traditional financial institutions, such as the financial sector's current dependence on BigTech's cloud-based computing and storage (think of Amazon's AWS), technological advantages such as artificial intelligence, machine learning, and APIs, and regulatory advantages (BigTech isn't burdened with Dodd-Frank, Basel capital restrictions, model risk regulations, and anti-money laundering program regulations).

But what are the differences between "BigTech" and "FinTech"? Both papers provide definitions for, and examples of, the two terms:

BigTech

FSB: "refers to large technology companies that expand into the direct provision of financial services or of products very similar to financial products"

BIS: "refers to large, existing companies whose primary activity is in the provision of digital services, rather than mainly in financial services … BigTech companies offer financial products only as one part of a much broader set of business lines."

Both the FSB and BIS have the same BigTech firms: Facebook, Amazon, Apple, Google, Alibaba, Tencent, Vodafone, among others.

FinTech

FSB: "technology-enabled innovation in financial services that could result in new business models, applications, processes or products with an associated material effect on the provision of 'financial services' … used to describe firms whose business model focuses on these innovations."

BIS: "refers to technology-enabled innovation in financial services with associated new business models, applications, processes, or products, all of which have a material effect on the provision of financial services."

Both the FSB and BIS use QuickenLoans and SOFI, among others, as examples of FinTech firms.

Which BigTech Firms are Providing What Financial Services Today?

The BIS paper provides a great summary table of the five main types of financial services that the eleven dominant BigTechs are currently providing. It's clear from this table that the three Chinese BigTechs – Alibaba, Tencent, and Baidu – have the most comprehensive suite of financial services/products, followed by the US trio of Google, Amazon, and Facebook. And BigTech is really big: The BIS paper notes that the six largest global BigTech firms all have market capitalizations greater than the market capitalization of the largest global financial institution, JPMorgan Chase.

Conclusion

There is no conclusion. Every day brings new entrants and participants, shifts, and changes. The regulatory environments are rapidly changing (although regulators and regulations always lag the regimes they regulate). But these two papers provide some insights into the world of FinTech, BigTech, and financial services, and are worth spending some time on.

FinCrime FinTech Hype, Hubris, and Subject Matter Enthusiasm

Posted December 20, 2018

This article came from my growing frustration with some of the new breed of technology experts/financial crimes enthusiasts.

Two very recent fincrime fintech start-ups recently published marketing papers – one a self-styled "Report" the other a blog – that should serve as reminders that, although innovation and change are critical to financial institutions' financial crimes risk management programs, fincrime fintechs are not. Or, put another way, those fincrime fintechs need to understand what they are and what they are not. Most important, they are not "solutions": they are tools that could be deployed, in whole or in part, by true financial crimes experts who bear the statutory and regulatory responsibility for – and personal liability for – designing, developing, implementing, maintaining, and enhancing their programs. And U.S. banking agencies are embracing the idea of responsibly implementing innovative approaches to financial crimes risk management. The U.S. banking agencies' December 3rd joint statement is a very positive step to encourage private sector innovation in fighting financial crime. But they don't limit those innovative approaches to just adopting new technologies: they also encourage "testing new ways of using existing tools". For those banks that are considering replacing their existing tools with "modern era technologies", I would caution them to first look at how they are using their existing tools, whether they have the data and in-house expertise to even deploy modern era technologies, and consider whether they are better off improving and augmenting their existing tools.

Let's take a look at the report and blog.

Feedzai – "the market leader in fighting financial crime fraud with AI"

The first report is from Feedzai, which, according to its website:

Feedzai is AI. We're coding the future of commerce with a leading platform powered by artificial intelligence and big data. Founded and developed by data scientists and aerospace engineers, Feedzai has one critical mission: make commerce safe. The world's largest banks, payment providers and retailers use Feedzai's machine learning technology to manage risks associated with banking and shopping, whether it's in person, online or via mobile devices.

[and …]

Feedzai is the market leader in fighting financial crime fraud with AI. But even a leader needs partners. To maximize our impact, we partner with top tier financial institutions, consultancies, system integrators, and technology providers to create win/win/win scenarios for the marketplace.

Feedzai's report is titled "A Guide for Financial Institutions – Augmenting Your AML with AI: See The Risk Signals in the Noise"

https://assets.sourcemedia.com/01/2a/181664d6497aae451c6911ebb6f4/feedzai-aml-report-v106.pdf

This "report" is really a marketing document from Feedzai, used to convince financial institutions that if they're not deploying machine learning and AI – indeed, if they're not deploying Feedzai's machine

learning and AI – they're at risk of what they refer to as "The Six Pains of Money Laundering" which can only be addressed if the buyer "flips the script with Feedzai anti-money laundering."

Let's look at those six pains. First, and foremost, none of them are actually pains *of money laundering*, but of complying with government-imposed legislative and regulatory requirements and expectations. Some aren't even pains, or pains related to the technology solutions that Feedzai is selling, but simple observations.

The first "pain" is ***regulatory fines***. Feedzai notes that "In the past decade, compliance fines erased $342 billion in profits for top US and European banks. This figure is expected to exceed $400 billion by 2020." And then they list what is implied to be AML-related fines for 11 banks and 1 non-bank telecom manufacturer. Going through those, *every one of them is solely, or primarily, an OFAC or sanctions-related penalty, with AML either not part of the penalty or, in the case of the hybrid OFAC/AML penalties, a small part.* At best Feedzai's list is sloppy and incomplete: at worst it is deceptive. If they're going to write a paper touting their AML capabilities that includes regulatory fines as the first pain point, they could at least use AML-related regulatory fines.

The second "pain" is ***organizational burden***. They write: "Financial institutions might employ upwards of 5,000 employees in sanction screening alone. As transaction volume keeps growing, so do alerts, false positives, and compliance teams, all at unsustainable rates." Again, they've confused AML with sanctions. Economic sanctions programs are related to AML programs, just as fraud programs are related to AML programs. But they are very different disciplines and require very different programs, technologies, staffing, and reporting. And a phrase such as "might employ upwards of 5,000" is weak (the word "might") and ambiguous (does "upwards of 5,000" mean 4,900? 1,000?).

The third "pain" is that "***current transaction monitoring solutions lack context***". Feedzai writes:

A PwC report states that transaction monitoring for AML often generates false positive rates of over 90%. The rule based systems that monitor these transactions do what they were supposed to: point to incidents where money movement exceeded certain thresholds. However, compliance teams cannot go deeper to provide additional context that would substantiate or refute the actual money laundering risk. Current solutions are unable to connect the dots between multiple seemingly unrelated alerts in order to contextualize and visualize suspicious movement patterns that point to broader AML risk.

First, the reason that there are false positives is that compliance teams *must, can, and do* go deeper than the alert generating monitoring systems to provide additional context to substantiate (apparently in 10% of the cases) or refute (in 90%). But those teams don't substantiate or refute *"the actual money laundering risk"* as Feedzai writes. What financial institutions are charged with is making a determination that certain activity is *suspicious*, not that it is, in fact, *money laundering*. And as all experienced AML professionals know, it is the job of the analyst or investigator to take the alert or referral and to determine whether the activity has no business or apparent lawful purposes or is not the type of activity that the particular customer would normally be expected to engage in, and to conclude that there is no reasonable explanation for the activity after examining the available facts, including the background and possible purpose of the transactions and activity. It is fair, though, that analysts and the entire financial services industry would be better served if AML transaction monitoring, interaction monitoring, and customer surveillance applications could produce alerts that led to SARS in more than 10% of the cases.

But as I will write in an upcoming article, addressing the false positive issue is more about, or at least as much about, cleaning up a bank's data and regulatory reform, than it is about deploying new technology.

Second, if a bank's current solution is "unable to connect the dots between multiple seemingly unrelated alerts in order to contextualize and visualize suspicious movement patterns that point to broader AML risk", then that bank is not using the data it has available to it in any reasonable way. A simple Scenario Analysis tool, such as the one I first developed in 1999 (and the subject of a July 2018 News post on this site), was used to run sophisticated, segmented customer surveillance models using basic relational database tools. That, coupled with a rudimentary case management system that allowed grouping and de-duplicating of related alerts and referrals into consolidated case packages, connected the dots in two different multi-national financial institutions. Connecting AML dots does not require banks to rip-and-replace existing tools: it requires them to creatively use their existing tools.

The fourth "pain of money laundering" that Feedzai identifies is ***manual SAR reporting***. But their description of this manual reporting pain point doesn't really address the manual nature of the process nor offer a technology solution. They write:

Typically as little as 7% of all filed SARS are deemed by the regulator as worthy of further AML investigation, which means that 95% of the effort of these teams goes to waste. As SAR reporting is still a highly manually intensive task, the end result is that most of the AML resources allocated by FIs and the regulator are busy clearing their own "noise," created in the first place because they are unable to substantiate true money laundering risk. Today's compliance-focused systems use limited legacy technologies and reward quantity over quality, sending millions of dollars to waste.

First, regulators (at least US regulators) don't examine banks on whether their SARs are "worthy of further AML investigation". It may be that the 7 per cent figure used by Feedzai reports to the largest banks anecdotal statements that they get some sort of law enforcement response to roughly 7% of their SARs, with responses being a follow up subpoena, a formal request for supporting documentation, or a national security letter. That doesn't mean that the other 93% of SARs "go to waste". I recently wrote that law enforcement (in the case of the FBI) can conservatively say that at least 20% of BSA filings have tactical or strategic value to law enforcement. We would all like to see that percentage go up, and that is a noble task for Feedzai, other fintechs, the financial sector, regulators, and law enforcement.

Second, I'm not sure what Feedzai means by writing that "most of the AML resources allocated by FIs and the regulator are busy clearing their own 'noise,' created in the first place because they are unable to substantiate true money laundering risk." Including regulators in this statement is confusing (to me) and it suggests that regulators are allocating resources (Q. their own resources or compelling banks to allocate bank resources) because regulators cannot substantiate true money laundering risk.

The fifth "pain" is ***disconnected business units***, and Feedzai impugns siloed AML and fraud units, and disconnected investigations and analytics teams. Both are, indeed, pain points for any program, but both are easily overcome without deploying any new technology. They are organizational problems overcome with organizational solutions.

The sixth and last "pain" is the "***barrier to digital transformation***." Feedzai describes this pain not as a barrier to digital transformation but *because* of digital transformation, because this digital transformation across the bank's businesses and operations "can harbor new waves of financial crime with criminals

hiding behind large new sets of distributed and disconnected data." The solution? "The magnitude of the detection complexity calls for new technologies to take the helm as legacy systems simply don't scale up to the task."

With these pains, Feedzai concludes that banks must "flip the script with Feedzai anti-money laundering". They announce "the dawn of machine learning for AML" with Feedzai's machine learning and advanced automation, etc.

Unfortunately, it's not the dawn of machine learning for AML. It may be the dawn for some banks that have allowed their programs and technologies to stagnate and become obsolete. But for five to ten years there have been banks (Wells Fargo) and fintechs (Verafin) using machine learning, artificial intelligence, and visual (geographical, temporal, relational) analytics to "replace the manually tedious parts of existing AML processes with insights that are specific to money laundering" to "separate meaningful risk signals from noise, ensuring that manual investigation resources are applied using a validated risk-based approach" and to allow FIU analysts to "understand suspicious patterns and more precisely allocate their manual investigation resources", all using advanced financial crimes-specific case management applications to ingest, triage, de-duplicate, risk-score, package, decision, and route alerts and referrals; triage, risk-score, and make SAR decisions; automate and write narratives; manage and report to external and internal stakeholders; and feed all of this back into the system to learn and adapt, tune and adapt models, and revise customer risk ratings.

So Feedzai: if you believe you are the best, or want to be the best, AML systems provider in the industry, your marketing materials such as "A Guide for Financial Institutions – Augmenting Your AML with AI: See The Risk Signals in the Noise" should be the best. They're not. Your subject matter enthusiasm is to be commended; your subject matter expertise needs work.

Tookitaki – intending to transform the way organizations do predictive modeling

According to its website …

Tookitaki is building an intelligent decision support system (DSS) to help businesses take smarter decisions. Built on an effective AI system, our DSS intends to transform the way organisations do predictive modeling. Most businesses globally use consultants, build ad hoc predictive models on sample data and take decisions. The current process offers neither efficiency nor scale – rather becomes obsolete in the world of big data.

Our DSS will empower businesses go beyond the barriers of existing statistical packages creating one-off solutions by offering production-ready, automated predictive modeling. Clients can call our REST API for live feedback and take actions accordingly.

Tookitaki's CEO, Abhishek Chatterjee, published a blog on December 14, 2018 titled "Modern Tech to Reshape US AML Compliance with Regulators' Recent Handshake." Let's take a look at that blog.

Mr. Chatterjee begins with his synopsis of the Joint Statement on Innovative Efforts to Combat Money Laundering and Terrorist Financing:

On December 3, The Board of Governors of the Federal Reserve System, the Federal Deposit Insurance Corporation (FDIC), the Financial Crimes Enforcement Network (FinCEN), the National Credit Union Administration, and the Office of the Comptroller of the Currency issued a joint statement encouraging banks to use modern-era technologies to bolster their Bank Secrecy Act/anti-money laundering (BSA/AML) compliance programs. The agencies ask banks "to consider, evaluate, and, where appropriate, responsibly implement innovative approaches to meet their Bank Secrecy Act/anti-money laundering (BSA/AML) compliance obligations, in order to further strengthen the financial system against illicit financial activity.

Actually, the agencies did not issue a statement encouraging banks to use modern-era technologies to bolster their BSA/AML programs. The agencies' statement encouraged banks to "consider, evaluate, and, where appropriate, responsibly implement *innovative approaches* to meet their" BSA/AML compliance obligations". And, in the very next sentence following the quote above, the Joint Statement provides, "[t]he Agencies recognize that private sector innovation, *including new ways of using existing tools* or adopting new technologies, can help banks …".

Notably, the Agencies are *not* limiting innovative approaches to the adoption of new ("modern-era") technologies (and by implication, replacement of not-so-modern-era technologies), but including new ways of using existing tools. This is critically important to those banks that are facing increasing pressure from fincrime fintechs to rip-and-replace existing AML systems with new, and often untested, technologies.

They are of the view that private sector innovation, involving new technologies such as artificial intelligence and machine learning, can help banks identify and report money laundering, terrorist financing and other illicit activities.

The Agencies provide two examples of innovative approaches: the use of innovative Financial Intelligence Units (FIUs) and "artificial intelligence and digital identity technologies". Notably, bank FIUs have been in existence since the late 1990s (I know, I deployed the first large bank FIU at FleetBoston Financial in 1999). The concept of a bank FIU is twenty years old, and almost every large financial institution now has an FIU that is continually implementing innovative approaches to fighting financial crimes. The success of an FIU is equal parts data, technology, tools, courage, imagination, compassion, empathy, cynicism, collaboration, hard work, patience, and luck.

Mr. Chatterjee next describes the "assurances" the agencies give:

In addition, the regulators assured that they will not penalize those firms who are found to have a deficiency in their existing compliance programs as they run pilots employing modern technologies. The statement reads: "While the Agencies may provide feedback, pilot programs in and of themselves should not subject banks to supervisory criticism even if the pilot programs ultimately prove unsuccessful. Likewise, pilot programs that expose gaps in a BSA/AML compliance program will not necessarily result in supervisory action with respect to that program." They have added that "the implementation of innovative approaches in banks' BSA/AML compliance programs will not result in additional regulatory expectations."

This is a reasonably accurate description of the assurances – although I would not use the word "assurances" given the qualifiers attached to it. The first three "assurances", and two more, are clear cut:

1. "The Agencies recognize that private sector innovation, including new ways of using existing tools or adopting new technologies, can help banks identify and report money laundering, terrorist financing, and other illicit financial activity by enhancing the effectiveness and efficiency of banks' BSA/AML compliance programs. To assist banks in this effort, *the Agencies are committed to continued engagement with the private sector and other interested parties."*
2. *"The Agencies will not penalize or criticize banks that maintain effective BSA/AML compliance programs commensurate with their risk profiles but choose not to pursue innovative approaches*."
3. "While banks are expected to maintain effective BSA/AML compliance programs, *the Agencies will not advocate a particular method or technology for banks to comply with BSA/AML requirements."*
4. Where test or implemented "artificial intelligence-based transaction monitoring systems … identify suspicious activity that would not otherwise have been identified under existing processes, *the Agencies will assess the adequacy of banks' existing suspicious activity monitoring processes independent of the results of the pilot program*"
5. "… the implementation of innovative approaches in banks' BSA/AML compliance programs *will not result in additional regulatory expectations.*"

Note the strong, unqualified language: "the Agencies *are committed* to continued engagement", "the Agencies *will not* penalize or criticize", "the Agencies *will not* advocate …", "the Agencies *will* assess", and "the implementation of innovative approaches *will not* result in additional regulatory expectations".

The qualified "assurances" come in the paragraph about pilot programs (with emphasis added):

Pilot programs undertaken by banks, in conjunction with existing BSA/AML processes, are an important means of testing and validating the effectiveness of innovative approaches. While the Agencies may provide feedback, pilot programs in and of themselves *should not* subject banks to supervisory criticism even if the pilot programs ultimately prove unsuccessful. Likewise, pilot programs that expose gaps in a BSA/AML compliance program *will not necessarily result in supervisory action* with respect to that program. For example, when banks test or implement artificial intelligence-based transaction monitoring systems and identify suspicious activity that would not otherwise have been identified under existing processes, the *Agencies will not automatically assume* that the banks' existing processes are deficient. In these instances, the Agencies will assess the adequacy of banks' existing suspicious activity monitoring processes independent of the results of the pilot program. Further, the implementation of innovative approaches in banks' BSA/AML compliance programs will not result in additional regulatory expectations.

Here there are the qualified assurances (which are not assurances): "should not", "will not necessarily", and "not automatically assume". These are important distinctions. The Agencies could have written something very different:

"… pilot programs in and of themselves *will not* subject banks to supervisory criticism even if the pilot programs ultimately prove unsuccessful. Likewise, pilot programs that expose gaps in a BSA/AML compliance program *will not result in supervisory action* with respect to that program. For example, when banks test or implement artificial intelligence-based transaction monitoring systems and identify suspicious activity that would not otherwise have been identified under existing processes, the *Agencies will not assume* that the banks' existing processes are deficient …"

But the author of the blog also uses an interesting qualifier by writing that the joint statement "largely clears the air for modern AML solutions, especially those based on artificial intelligence and machine learning". I agree: the joint statement largely, but not entirely, clears the air or provides some comfort to banks who implement innovative approaches, including machine learning and AI. But as the Agencies remind us, any innovative approaches must be done responsibly while the bank continues to meet its BSA/AML program obligations and, if in doing so any gaps in that program that are identified will not necessarily result in supervisory action, but the Agency will assess those gaps to determine whether the program is, in fact, meeting regulatory requirements.

Finally, I disagree with Mr. Chatterjee's statement that we are in an "era of sophisticated financial crimes that are impossible to detect with legacy systems." I trust that this is simply a marketing phrase, and the use of the absolute word "impossible" is puffery and salesmanship. The statement is false.

Like Mr. Chatterjee and his firm, I also am "both happy and excited at the US regulators' change of tone with regard to the use of modern technologies by banks and financial institutions to combat financial crimes such as money laundering." But we need to be as realistic and practical as we are happy and excited about embracing new technologies without fully utilizing the existing technologies. Modern era technologies will be no better than the existing technologies if they are deployed against incomplete, outdated, stale, poorly labeled data by people lacking courage, imagination, and financial crimes expertise.

The U.S. banking agencies' December 3rd joint statement is a very positive step to encourage private sector innovation in fighting financial crime by testing new ways of using existing tools as well as adopting new technologies. For those banks that are considering replacing their existing tools with "modern era technologies", I would caution them to first look at how they are using their existing tools, whether they have the data and in-house expertise to even deploy modern era technologies, and consider whether they are better off improving and augmenting their existing tools. A bank's data and personnel are the "rails" upon which the AML technology rides: if those rails can't support the high-speed train of machine learning and AI-based systems, then it's best to fix and replace the rails before you test and buy the new train.

Regulatory Lag & Drag – Are There FinTech Solutions?

The RegTech, SupTech, and FinTech communities are focused on developing new technologies to speed up, simplify, and streamline financial institutions' ability to implement new rules, regulations, and regulatory guidance. But there are two other stages of the regulatory life cycle that may be longer and more problematic for financial institutions than implementing new regulations: these are the time it takes for new regulations to be written and published ("Regulatory Lag"), and the time it takes to enforce those regulations ("Regulatory Drag").

> Posted January 28, 2019
>
> Fintech/Regtech proponents believe that technology can solve most of the inefficiencies in the current regulatory environment. This article points out two areas where technology may not have a solution …

Time to Regulate – or "Regulatory Lag"

This lag occurs where a new risk emerges, or a new product is introduced, or an existing product is used in new ways. There is always a lag between that new risk or product and the resulting legislative and/or regulatory response. In the meantime, institutions have to begin addressing the new risks when they first

emerge – they can't wait for new rules, regulatory guidance, and regulations to begin the multi-year people, process, and technology changes necessary to address the requirements of the regulation. Those early, pre-rule and pre-regulation efforts at building controls to address new risks can be expensive, and institutions run the risk of missing the mark and having to re-do much of what they've built. The best example of regulatory lag in the AML space is 9/11, which saw legislation passed in 45 days (October 2001), regulations published two years later (2003), and regulatory guidance in the form of the BSA Exam Manual two years after that (2005). Although it was only 45 days that financial institutions knew about the new information sharing provisions in section 314 of the USA PATRIOT Act, it was almost another four years before financial institutions knew how their regulators would examine their compliance with those information sharing provisions. It was this "regulatory lag" that led to my written statement (in December 2006) that "we'll be judged tomorrow on what we're building today, based on regulations that haven't yet been written and best practices that haven't been shared."

Time to Enforce – or "Regulatory Drag"

Public enforcement actions (and prosecutions) drive a lot of compliance-related behavior in financial services. Yet there are multi-year delays between when the impugned behavior occurred and when a public enforcement action (and/or prosecution) makes them known to the industry. FinCEN's December 2014 action against MoneyGram's former BSA Officer is a good example: that action was made public in December 2014, and alleged violations of the Bank Secrecy Act that occurred from 2003 through May 2008, or more than 6 ½ years from the last day of the impugned activity and when the public action was taken.

What Can Technology Do To Address Regulatory Lag and Drag?

Regulatory lag and drag have been around for as long as there have been regulators. But with the world speeding up as much as it is, with new products and services, and new providers, being rolled out and created much faster than regulatory bodies can manage, there must be changes made in the entire regulatory life cycle.

FinTech providers and their customers demand a fast revolution. Regulators prefer a slow, deliberate evolution. There has to be a better way to identify new and emerging risks, to draft and communicate regulations to address those risks, and to implement the needed controls to manage those risks.

I'm not sure what can be done from a purely technology perspective to speed up regulators (and prosecutors), but the proponents of FinTech, RegTech, and SupTech solutions shouldn't just focus on digitizing the implementation of new regulations, but on digitizing the entire regulatory life cycle: the regulatory lag between new risks and new regulations, the regulations themselves, and the regulatory drag from regulatory problem to public resolution.

Rules-Based Monitoring, Alert to SAR Ratios, and False Positive Rates – Are We Having The Right Conversations?

Posted December 20, 2018

There is a lot of conversation in the industry about the inefficiencies of "traditional" rules-based monitoring systems, Alert-to-SAR ratios, and the problem of high false positive rates. Let me add to that conversation by throwing out what could be some controversial observations and suggestions …

Current Rules-Based Transaction Monitoring Systems – are they really that inefficient?

For the last few years AML experts have been stating that rules-based or typology-driven transaction monitoring strategies that have been deployed for the last 20 years are not effective, with high false positive rates (95% false positives!) and enormous staffing costs to review and disposition all of the alerts. Should these statements be challenged? Is it the fact the transaction monitoring strategies are rules-based or typology-driven that drives inefficiencies, or is it the fear of missing something driving the tuning of those strategies? Put another way, if we tuned those strategies so that they only produced SARs that law enforcement was interested in, we wouldn't have high false positive rates and high staffing costs. Graham Bailey, Global Head of Financial Crimes Analytics at Wells Fargo, believes it is a combination of basic rules-based strategies coupled with the fear of missing a case. He writes that some banks have created their staffing and cost problems by failing to tune their strategies, and by "throwing orders of magnitude higher resources at their alerting." He notes that this has a "double negative impact" because "you then have so many bad alerts in some banks that they then run into investigators' 'repetition bias', where an investigator has had so many bad alerts that they assume the next one is already bad" and they don't file a SAR. So not only are the SAR/alert rates so low, you run the risk of missing the good cases.

After 20+ years in the AML/CTF field – designing, building, running, tuning, and revising programs in multiple global banks – I am convinced that rules-based interaction monitoring and customer surveillance systems, running against all of the data and information available to a financial institution, managed and tuned by innovative, creative, courageous financial crimes subject matter experts, can result in an effective, efficient, proactive program that both provides timely, actionable intelligence to law enforcement and meets and exceeds all regulatory obligations. Can cloud-based, cross-institutional, machine learning-based technologies assist in those efforts? Yes! If properly deployed and if running against all of the data and information available to a financial institution, managed and tuned by innovative, creative, courageous financial crimes subject matter experts.

Alert to SAR Ratios – is that a ratio that we should be focused on?

A recent Mid-Size Bank Coalition of America (MBCA) survey found the average MBCA bank had: 9,648,000 transactions/month being monitored, resulting in 3,908 alerts/month (0.04% of transactions alerted), resulting in 348 cases being opened (8.9% of alerts became a case), resulting in 108 SARs being filed (31% of cases or 2.8% of alerts). Note that the survey didn't ask whether any of those SARs were of interest or useful to law enforcement. Some of the mega banks indicate that law enforcement shows interest in (through requests for supporting documentation or grand jury subpoenas) 6% – 8% of SARs.

So I argue that the Alert/SAR and even Case/SAR (in the case of Wells, Package/Case and Package/SAR) ratios are all of interest, but tracking to SARs filed is a little bit like a car manufacturer tracking how many cars it builds but not how many cars it sells, or how well those cars perform, how well they last, and how popular they are. The better measure for AML programs is "SARs purchased", or SARs that provide value to law enforcement.

How do you determine whether a SAR provides value to Law Enforcement? One way would be to ask Law Enforcement, and hope you get an answer. That could prove to be difficult. Can you somehow measure Law Enforcement interest in a SAR? Many banks do that by tracking grand jury subpoenas received to prior SAR suspects, Law Enforcement requests for supporting documentation, and other formal and informal requests for SARs and SAR-related information. As I write above, an Alert-to-SAR rate may not be

a good measure of whether an alert is, in fact, "positive". What may be relevant is an Alert-to-TSV SAR rate (see my previous article for more detail on TSV SARs). What is a "TSV SAR"? A SAR that has Tactical or Strategic Value to Law Enforcement, where the value is determined by Law Enforcement providing a response or feedback to the filing financial institution within five years of the filing of the SAR that the SAR provided tactical (it led to or supported a particular case) or strategic (it contributed to or confirmed a typology) value. If the filing financial institution does not receive a TSV SAR response or feedback from law enforcement or FinCEN within five years of filing a SAR, it can conclude that the SAR had no tactical or strategic value to law enforcement or FinCEN, and may factor that into decisions whether to change or maintain the underlying alerting methodology. Over time, the financial institution could eliminate those alerts that were not providing timely, actionable intelligence to law enforcement, and when that information is shared across the industry, others could also reduce their false positive rates.

Which leads to …

False Positive Rates – if 95% is bad … what's good?

There is a lot of lamenting, and a lot of axiomatic statements, about high false positive rates for AML alerts: 95% or even 98% false positive rates. I'd make three points.

First, vendors selling their latest products, touting machine learning and artificial intelligence as the solution to high false positive rates, are doing what they should be doing: convincing consumers that their current product is out-dated and ill-equipped for its purpose by touting the next, new product. I argue that high false positive rates are not caused by the current rules-based technologies; rather, they're caused by inexperienced AML enthusiasts or overwhelmed AML experts applying rules that are too simple against data that is mis-labeled, incomplete, or simply wrong, and erring on the side of over-alerting and over-filing for fear of regulatory criticism and sanctions.

If the regulatory problems with AML transaction monitoring were truly technology problems, then the technology providers would be sanctioned by the regulators and prosecutors. But an AML technology provider has never been publicly sanctioned by regulators or prosecutors … for the simple reason that any issues with AML technology aren't technology issues: they are operator issues.

Second, are these actually "false" alerts? Rather, they are alerts that, at the present time, based on the information currently available, do not rise to the level of either (i) requiring a complete investigation, or (ii) if completely investigated, do not meet the definition of "suspicious". Regardless, they are now valuable data points that go back into your monitoring and case systems and are "hibernated" and possibly come back if that account or customer alerts at a later time, or there is another internally- or externally-generated reason to investigate that account or customer.

Third, if 95% or 98% false positive rates are bad … what is good? What should the target rate be? I'll provide some guidance, taken from a Treasury Office of Inspector General (OIG) Report: OIG-17-055 issued September 18, 2017 titled "FinCEN's information sharing programs are useful but need FinCEN's attention." The OIG looked at 314(a) statistics for three years (fiscal years 2010-2012) and found that there were 711 314(a) requests naming 8,500 subjects of interest sent out by FinCEN to 22,000 financial institutions. Those requests came from 43 Law Enforcement Agencies (LEAs), with 79% of them coming from just six LEAs (DEA, FBI, ICE, IRS-CI, USSS, and US Attorneys' offices). Those 711 requests resulted in 50,000 "hits" against customer or transaction records by 2,400 financial institutions.

To analogize those 314(a) requests and responses to monitoring alerts, there were 2,400 "alerts" (financial institutions with positive matches) out of 22,000 "transactions" (total financial institutions receiving the 314(a) requests). That is an 11% hit rate or, arguably, a 89% false positive rate. And keep in mind that in order to be included in a 314(a) request, the Law Enforcement Agency must certify to FinCEN that the target "is engaged in, or is reasonably suspected based on credible evidence of engaging in, terrorist activity or money laundering." So Law Enforcement considered that all 8,500 of the targets in the 711 requests were active terrorists or money launderers, and 11% of the financial institutions positively responded.

With that, one could argue that a "hit rate" of 10% to 15% could be optimal for any reasonably designed, reasonably effective AML monitoring application.

But a better target rate for machine-generated alerts is the rate generated by humans. Bank employees – whether bank tellers, relationship managers, or back-office personnel – all have the regulatory obligation of reporting unusual activity or transactions to the internal bank team that is responsible for managing the AML program and filing SARs. For the twenty plus years I was a BSA Officer or head of investigations at large multi-national US financial institutions, I found that those human-generated referrals resulted in a SAR roughly 40% to 50% of the time.

An alert to SAR ratio goal of machine-based alert generation systems should be to get to the 40% to 50% referral-to-SAR ratio of human-based referral generation programs.

Flipping the Three AML Ratios with Machine Learning and Artificial Intelligence (why Bartenders and AML Analysts will survive the AI Apocalypse)

Posted December 14, 2018

This is the most viewed article on my website.

Machine Learning and Artificial Intelligence proponents are convinced – and spend a lot of time trying to convince others – that they will disrupt and revolutionize the current "broken" AML regime. Among other targets within this broken regime is AML alert generation and disposition and reducing the false positive rate (more on false positives in another article!). The result, if we believe the ML/AI community, is a massive reduction in the number of AML analysts that are churning through the hundreds and thousands of alerts, looking for the very few that are "true positives" worthy of being labelled "suspicious" and reported to the government.

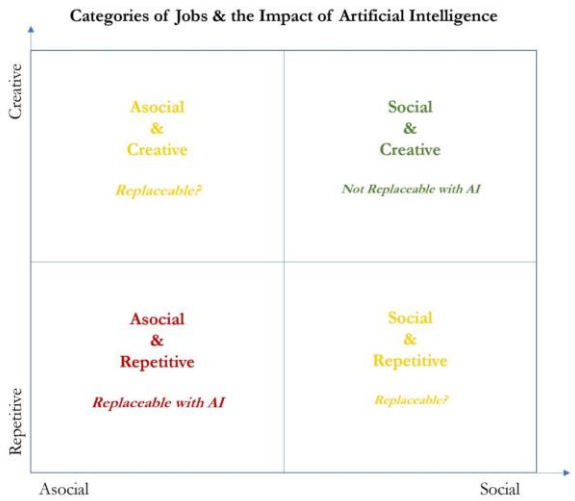

But is it that simple? Can the job of AML Analyst be eliminated or dramatically changed – in scope and number of positions – by machine learning and AI? Much has been and continues to be written about the impact of artificial intelligence on jobs.

Those writers have categorized jobs along two axes – a Repetitive-to-Creative axis, and an Asocial-to-Social axis – resulting in four "buckets" of jobs, with each bucket of jobs being more or less likely to be disrupted or even eliminated:

A good example is the "Social & Repetitive" job of Bartender: Bartenders spend much of their time doing very routine, repetitive tasks: after taking a drink order, they assemble the correct ingredients in the correct amounts, and put those ingredients in the correct glass, then present the drink to the customer. All of that could be more efficiently and effectively done with an AI-driven machine, with no spillage, no waste, and perfectly poured drinks. So why haven't we replaced bartenders? Because a good bartender has empathy, compassion, and instinct, and with experience can make sound judgments on what to pour a little differently, when to cut-off a customer, when to take more time or less with a customer. A good bartender adds value that a machine simply can't.

Another example could be the "Asocial & Creative" (or is it "Social & Repetitive"?) job of an AML Analyst: much of an AML Analyst's time is spent doing very routine, repetitive tasks: reviewing the alert, assembling the data and information needed to determine whether the activity is suspicious, writing the narrative. So why haven't we replaced AML Analysts? Because a good Analyst, like a good bartender, has empathy, compassion, and instinct, and with experience can make sound judgments on what to investigate a little differently, when to cut-off an investigation, when to take more time or less on an investigation. A good Analyst adds value that a machine simply can't.

Where AI and Machine Learning, and Robot Process Automation, can really help is by flipping the three currently inefficient AML ratios:

1. The False Positive Ratio – the currently accepted, but highly axiomatic and anecdotal, ratio is that 95% to 98% of alerts do not result in SARs, or are "false positives" … although no one has ever boldly stated what an *effective* or *acceptable* false positive rate is (even with ROC curves providing some empirical assistance), perhaps the ML/AI/RPA communities can flip this ratio so that 95% of alerts result in SARs. If they can do this, they can also convince the regulatory community that this new ratio meets regulatory expectations (because as I'll explain in an upcoming article, the false positive ratio problem may be more of a regulatory problem than a technology problem).
2. The Forgotten SAR Ratio – like false positive rates, there are anecdotes and some evidence that very few SARs provide tactical or strategic value to law enforcement. Recent Congressional testimony suggests that ~20% of SARs provide TSV (tactical or strategic value) to law enforcement … perhaps the ML/AI/RPA communities can help to flip this ratio so that 80% of SARs are TSV SARs. This also will take some effort from the regulatory and law enforcement communities.
3. The Analysts' Time Ratio – 90% of an AML Analyst's time can be spent simply assembling the data, information, and documents needed to investigate a case, and only 10% of their time thinking and using their empathy, compassion, instinct, judgment, and experience to make good decisions and file TSV SARs … perhaps the ML/AI/RPA communities can help to flip this ratio so that Analysts spend 10% of their time assembling and 90% of their time thinking.

We've seen great strides in the AML world in the last 5-10 years when it comes to applying machine learning and creative analytics to the problems of AML monitoring, alerting, triaging, packaging, investigations, and reporting. My good friend and former colleague Graham Bailey at Wells Fargo designed and deployed ML and AI systems for AML as far back as 2008-2009, and the folks at Verafin have deployed cloud-based machine learning tools and techniques to over 1,600 banks and credit unions.

I've outlined three rather audacious goals for the machine learning/artificial intelligence/robotic process automation communities:

1. The False Positive Ratio – flip it from 95% false positives to 5% false positives
2. The Forgotten SAR Ratio – flip it from 20% TSV SARs to 80% TSV SARs
3. The Analysts' Time Ratio – flip it from 90% gathering data to 10% gathering data

Although many new AML-related jobs are being added – data scientist, model validator, etc. – and many existing AML-related jobs are changing, I am convinced that the job of AML Analyst will always be required. Hopefully, it will shift over time from being predominantly that of a gatherer of information and more of a hunter of criminals and terrorists. But it will always exist. If not, I can always fall back on being a Bartender. Maybe …