

May 2022

The case for placing AI at the heart of digitally robust financial regulation

Jo Ann Barefoot

This report is available online at: <https://www.brookings.edu/center/center-on-regulation-and-markets/>

B | Center on
Regulation and Markets
at BROOKINGS

The Center on Regulation and Markets at Brookings provides independent, non-partisan research on regulatory policy, applied broadly across microeconomic fields. It creates and promotes independent economic scholarship to inform regulatory policymaking, the regulatory process, and the efficient and equitable functioning of economic markets.



Affiliations

Jo Ann Barefoot, CEO & Cofounder, Alliance for Innovative Regulation

Disclosure

The author serves on the board of directors of FinRegLab, a nonprofit organization whose research includes a focus on use of AI in financial regulatory matters. She did not receive financial support from any firm or person for this article or from any firm or person with a financial or political interest in this article. Other than the aforementioned, the author is not currently an officer, director, or board member of any organization with a financial or political interest in this article.

Introduction

“Data is the new oil.”

[Originally coined](#) in 2006 by the British mathematician Clive Humby, this phrase is arguably more apt today than it was then, as smartphones rival automobiles for relevance and the technology giants know more about us than we would like to admit.

Just as it does for the financial services industry, the hyper-digitization of the economy presents both opportunity and potential peril for financial regulators. On the upside, reams of information are newly within their reach, filled with signals about financial system risks that regulators spend their days trying to understand. The explosion of data sheds light on global money movement, economic trends, customer onboarding decisions, quality of loan underwriting, noncompliance with regulations, financial institutions' efforts to reach the underserved, and much more. Importantly, it also contains the answers to regulators' questions about the risks of new technology itself. Digitization of finance generates novel kinds of hazards and accelerates their development. Problems can flare up between scheduled regulatory examinations and can accumulate imperceptibly beneath the surface of information reflected in traditional reports. Thanks to digitization, regulators today have a chance to gather and analyze much more data and to see much of it in something close to real time.

The potential for peril arises from the concern that the regulators' current technology framework lacks the capacity to synthesize the data. The irony is that this flood of information is too much for them to handle. Without digital improvements, the data fuel that financial regulators need to supervise the system will merely make them overheat.

Enter artificial intelligence.

In 2019, then-Bank of England Gov. Mark Carney argued that financial regulators will have to adopt AI techniques in order to keep up with the rising volumes of data flowing into their systems. To dramatize the point, he [said the bank](#) receives 65 billion pieces of data annually from companies it oversees and that reviewing it all would be like “each supervisor reading the complete works of Shakespeare twice a week, every week of the year.”

That was three years ago. The number is almost certainly higher today. Furthermore, the numbers he cited only covered information reported by regulated firms. It omitted the massive volumes of external “Big Data” generated from other sources like public records, news media, and social media that regulators should also be mining for insight about risks and other trends.

AI was developed [over 70 years ago](#). For decades, enthusiasts predicted that it would change our lives profoundly, but it took awhile before AI had much impact on everyday lives.¹ AI occasionally made news by performing clever feats, like IBM's Watson besting [human champions at Jeopardy](#) in 2011, or AIs beating masters of complex games like [chess](#) (in 1996) and [Go \(in 2017\)](#). However, it was only recently that such machines showed signs of being able to solve real-world problems. Why is that?

A key answer is that, until only recently, there wasn't enough data in digitized form—formatted as computer-readable code—to justify using AI.² Today, there is so much data that not only *can* we use AI, but in many fields like financial regulation we *have* to use AI simply to keep up.

¹ Anyoha, Rockwell. “The History of Artificial Intelligence,” Harvard University, August 8, 2017. <https://sitn.hms.harvard.edu/flash/2017/history-artificial-intelligence/>.

² McKendrick, Joe. “AI Adoption Skyrocketed Over the Last 18 Months.” Harvard Business Review, September 27, 2021. <https://hbr.org/2021/09/ai-adoption-skyrocketed-over-the-last-18-months;> McKendrick, Joe. “The Data Paradox: Artificial Intelligence Needs Data; Data Needs AI.” Forbes,

As discussed further below, financial regulators around the world are in the early stages of exploring how AI and its sub-branches of Machine Learning (ML), Natural Language Processing (NLP), and neural networks, can enhance their work. They are increasingly weighing the adoption of “supervisory technology” (or “suptech”) to monitor companies more efficiently than they can with analog tools. This shift is being mirrored in the financial industry by a move to improve compliance systems with similar “regulatory technology” (“regtech”) techniques. Both processes are running on a dual track, with one goal being to convert data into a digitized form and the other to analyze it algorithmically. Meeting either of these objectives without the other has little value. Together, they will transform both financial regulation and compliance. They offer the promise that regulation, like everything else that gets digitized, can become better, cheaper, and faster, all at once.

Imagine if AI was already the default mechanism

Financial regulators around the world have generally been more active in [regulating industry’s use of AI](#) than adopting it for their own benefit. Opportunities abound, however, for AI-powered regulatory and law enforcement tactics to combat real-world problems in the financial system. In a later section, this paper will look at the primary emerging use cases. Before doing so, it is worth taking a look at some areas of poor regulatory performance, both past and present, and ask whether AI could have done better.

One example is the \$800 billion Paycheck Protection Program that Congress established in 2020 to provide government-backed loans for small businesses reeling from the pandemic. More than 15% of PPP “loans” —representing \$76 billion—contained [evidence of fraud](#), according to a study released last year. Many cases involved loan applicants using fake identities. Imagine if the lenders submitting loan guarantee applications or the Small Business Administration systems that were reviewing them had had mature AI-based systems that could have flagged suspicious behavior. They could have spotted false statements and prevented fraudulent loans, thereby protecting taxpayer money and ensuring that their precious funds helped small businesses in need instead of financing thieves.

Two examples can be found from the war in Ukraine. The Russian invasion has sparked a [whole new array of sanctions](#) against Russian oligarchs who hide riches in shell companies and are scrambling to move their money undetected. Financial institutions are required to screen accounts and transactions to identify transactions by sanctioned entities. What if they and law enforcement agencies like the Financial Crimes Enforcement Network (FinCEN) had AI-powered analytics to pull and pool data from across the spectrum of global transactions and find the patterns revealing activity by sanctioned parties? Unfortunately, most financial institutions and government agencies do not have these tools in hand today.

The second example comes from the rapid flight of millions of refugees [attracting human traffickers](#) to the country’s borders seeking to ensnare desperate women and children and sell them into slavery for work and sex. Banks are required by law to maintain anti-money laundering (AML) systems to detect and report money movement that may indicate human trafficking and other crimes, but these systems are mostly analog and [notoriously ineffective](#). The United Nations’ Office on Drugs and Crime estimates that [less than 1% of financial crime](#) is caught. AI-powered compliance systems would have a far better chance of flagging the criminal rings targeting Ukraine. If such systems had been in effect in recent years, moreover, the human trafficking trade might not be flourishing. As it stands today, an estimated 40 million people are being held [captive in modern human slavery](#), and one in four of them is a child.

In another thought experiment, what if bank regulators in 2007 had been able to see the full extent of interrelationships between subprime mortgage lenders and Wall Street firms like Bear Stearns, Lehman Brothers, and AIG? If regulators had been armed with real-time digital data and AI analytics, they would

September 27, 2021. <https://www.forbes.com/sites/joemckendrick/2021/06/27/the-data-paradox-artificial-intelligence-needs-data-data-needs-ai/?sh=6cd30dbe71a5>.

have been monitoring risk contagion in real time. They might have been able to avert the financial crisis and with it, the Great Recession.

Finally, what about fair lending? In 1968, the United States outlawed discrimination on the basis of race, religion and other factors in mortgage lending through the passage of the Fair Housing Act.³ With the later passage of the Equal Credit Opportunity Act and Housing and Community Development Act, both in 1974, Congress added sex discrimination to that list and expanded fair-lending enforcement to all types of credit, not just mortgages.⁴ That was nearly 50 years ago.

These laws have gone a long way toward combating straightforward, overt discrimination but have been much less effective in rooting out other forms of bias. Lending decisions still produce “disparate impacts” on different groups of borrowers, usually in ways that disproportionately harm protected classes like [people of color](#). Some of this arises from the fact that high volume credit decisioning must rely on efficient measures of creditworthiness, like credit scores, that in turn rely on narrow sources of data.⁵ What if, 40 years ago, both regulators and industry had been able to gather much more risk data and analyze it with AI? How many more people would have been deemed creditworthy instead of having their loan denied? Over four decades, could AI tools have changed the trajectory of racial opportunity in the United States, which currently includes a [\\$10 trillion racial wealth gap](#) and the African-American homeownership rate lagging that of whites [by 30 percentage points](#)?

How regulators aim to keep pace with exponentially changing tech

In his 2018 [book titled “Unscaled,”](#) venture capitalist Hemant Taneja argued that exploding amounts of data and AI will continue to produce unprecedented acceleration of our digital reality. “In another ten years anything that AI doesn’t power will seem lifeless and outmoded. It will be like an icebox after electric-powered refrigerators were invented,” he wrote.

Taneja’s estimated time horizon is now only six years away. In the financial sector, this sets up a daunting challenge for regulators to design and construct sufficiently powerful supotech before the industry’s changing technology could overwhelm their supervisory capacity. Fortunately, regulators in the U.S. and around the world are taking steps to narrow the gap.

Arguably the global leader in regulatory innovation is the United Kingdom’s Financial Conduct Authority (FCA). In 2015, the [FCA established the Project Innovate initiative](#), which included the creation of a “regulatory sandbox” for private sector firms to test new products for their regulatory impact. A year later, the FCA launched a regtech unit that developed what the agency called “techsprints”—an open competition resembling a tech hackathon in which regulatory, industry, and issue experts work side-by-side with software engineers and designers to develop and present tech prototypes for solving a particular regulatory problem. The innovation program has since been expanded into a major division within the FCA.⁶

The FCA has been able to translate this relatively early focus on digital innovation into real-world problem solving. In 2020, a senior agency official gave a [speech about how the FCA uses machine learning and natural language processing to monitor company behaviors and “spot outlier firms”](#) as part of a “holistic”

³ Federal Reserve. “Fair Housing Act summary.” Accessed May 20, 2022.

https://www.federalreserve.gov/boarddocs/supmanual/cch/fair_lend_fhact.pdf.

⁴ A Voice. “Fair Housing Timeline.” Accessed May 20, 2022. <http://www.avoiconline.org/fair-housing/timeline.html>.

⁵ Klein, Aaron. “Reducing bias in AI-based financial services.” Brookings Institution, July 10, 2020. <https://www.brookings.edu/research/reducing-bias-in-ai-based-financial-services/>.

⁶ Financial Conduct Authority. “FCA Innovation Hub.” Accessed May 20, 2022. <https://www.fca.org.uk/firms/innovation>.

approach to data analysis. Similar strides have been made in other countries, including [Singapore](#) and Australia.

U.S. regulators for the most part have made slower progress incorporating AI technologies in their monitoring of financial firms. All of the federal financial regulatory bodies have innovation programs in some form. Most of them, however, have focused more on industry innovation than their own. The U.S. banking agencies—[Consumer Financial Protection Bureau](#), [Federal Deposit Insurance Corporation](#), [Federal Reserve Board](#), and [Office of the Comptroller of the Currency](#)—all have innovation initiatives that are largely outward-facing, aimed at understanding new bank technologies and offering a point of contact on novel regulatory questions. They all also expanded their technology activities during the COVID-19 pandemic, spurred by the sudden digital shifts underway in the industry and their own need to expand offsite monitoring. Several agencies also have supotech projects underway. These, however, generally have limited reach and do not address the need for agencies to revisit their foundational, analog-era information architecture.

This is beginning to change. The Federal Reserve in 2021 created the new position of Chief Innovation Officer and hired Sunayna Tuteja from the private sector, charging her to undertake a sweeping modernization of the Fed’s data infrastructure. The FDIC, too, has closely examined its own data structures, and the OCC has worked on consolidating its examination platforms. These are productive steps, but they still lag the advanced thinking underway in other parts of the world. U.S. regulators have yet to narrow the gap between the accelerating innovation in the private sector and their own monitoring systems.

Other U.S. regulatory agencies have embraced AI technologies more quickly. In 2017, Scott Bauguess, the former deputy chief economist at the Securities and Exchange Commission (SEC), described his [agency’s use of AI](#) to monitor securities markets. Soon after the financial crisis, he said, the SEC began “simple text analytic methods” to determine if the agency could have predicted risks stemming from credit default swaps before the crisis. SEC staff also applies machine-learning algorithms to identify reporting outliers in regulatory filings.

Similarly, the Financial Industry Regulatory Authority (FINRA)—the self-regulatory body overseeing broker-dealers in the U.S.—uses robust AI to detect possible misconduct.⁷ The Commodity Futures Trading Commission (CFTC), meanwhile, has been a leader through its LabCFTC program, which addresses both fintech and regtech solutions. Former CFTC Chairman Christopher Giancarlo has said that the top priority of every regulatory body should be to “digitize the rulebook.”⁸ Lastly, the Treasury Department’s Financial Crimes Enforcement Network (FinCEN) launched an innovation program in 2019 to explore regtech methods for improving money-laundering detection.⁹ The agency is now in the process of implementing sweeping technology mandates it received under the Anti-Money Laundering Act of 2020, a great opportunity to implement AI to better detect some of the financial crimes discussed above.

Key financial regulation use cases

If government agencies supplanted their analog systems with a digitally native design, it would optimize the analysis of data that is now being under-utilized. The needles could be found in the haystack, fraudsters

⁷ Finra. “How the Cloud has Revolutionized FINRA Technology.” July 30, 2018. <https://www.finra.org/media-center/finra-unscripted/how-cloud-has-revolutionized-finra-technology>.

⁸ J. Christopher Giancarlo. “Quantitative Regulation: Effective Market Regulation in a Digital Era.” Commodity Futures Trading Commission. <https://www.cftc.gov/PressRoom/SpeechesTestimony/opagiancarlo59>.

⁹ Financial Crimes Enforcement Network. “Innovation Hours Program: Emerging Themes and Future Role in AML Act Implementation.” May 2019–February 2021. <https://www.fincen.gov/sites/default/files/2021-03/FinCEN%20IH%20Prgm%20Public%20Report%20508C.pdf>.

and money launderers would have a harder time hiding their activity, and regulators would more completely fulfill their mission of maintaining a safer and fairer financial system.

Below are specific use cases for incorporating AI in the regulatory process:

AML and sanctions screening

Arguably the most advanced regtech use case globally is anti-money laundering (AML). AML compliance costs the industry upwards of \$50 billion per year in the U.S., as most banks rely on rules-based transaction monitoring systems.¹⁰ These methods help them determine which activity to report to FinCEN as suspicious but currently produce a [false-positive rate](#) of over 90%. This suggests banks, regulators, and law enforcement authorities are spending time and money chasing down potential leads but not really curbing illicit financial crimes. The AML data that law enforcement agencies currently receive contains too much unimportant information and is not stored in formats to help identify patterns of crime.¹¹

In addition to the challenges associated with locating financial crimes among the massively complex web of global transactions, banks also must perform identity verification checks on new customers and submit “beneficial owner” data to FinCEN to prevent launderers from hiding behind fake shell companies. The war in Ukraine and toughening of sanctions on Russian oligarchs has highlighted the need for better screening mechanisms to restrict the financial activity of individuals that appear on sanctions lists. While a growing industry of regtech firms are attempting to help financial institutions more efficiently comply with Know-Your-Customer (KYC) rules, FinCEN is in the midst of implementing legislative reforms requiring corporations to submit data to a new beneficial owner database.

In 2018 and 2019, the FCA held two international tech sprints aimed at addressing AML challenges. The first sprint dealt with enabling regulators and law enforcement to share threat information more safely and effectively. The second focused on “Privacy-Enhancing Technologies,” or PET’s, of various kinds. For example, homomorphic encryption is a technique that shows promise for enabling data shared through AML processes to be encrypted throughout the analytical process, so that the underlying information is concealed from other parties and privacy is preserved. Another PET technique known as zero-knowledge proof enables one party to ask another essentially a yes-or-no question without the need to share the underlying details that spurred the inquiry. For example, one bank could ask another if a certain person is a customer, or if that person engaged in a certain transaction. Techniques like this can be used to enable machine-learning analysis of laundering patterns without compromising privacy or potentially undermining the secrecy of an ongoing investigation.

Fraud prevention

The SBA did make efforts to evaluate AI tools to detect fraud in PPP loans, looking to certain [AI-powered fintech lenders](#). Nevertheless, the small business loan program was still rife with fraud. (In fact, some of the attention regarding fraud concerns has centered on loans processed by fintech firms.¹²) Several [studies](#) show that effective use of machine learning in credit decisioning can more easily detect when, for example, loan applications are submitted by fake entities.

One of the biggest fraud threats facing financial institutions is the use of synthetic identities by bad actors. These are created by combining real customer information with fake data in a series of steps that can fool

¹⁰ LexisNexis Risk Solutions. “Create a Digitally-Driven Financial Crime Compliance Strategy” 2021. <https://risk.lexisnexis.com/insights-resources/research/true-cost-of-financial-crime-compliance-study-for-the-united-states-and-canada>.

¹¹ Nicodemus, Aaron. “Analysis: Improving SARs effectiveness takes more than act of Congress.” Compliance Week, January 19, 2021. <https://www.complianceweek.com/aml/analysis-improving-sars-effectiveness-takes-more-than-act-of-congress/29948.article>.

¹² Chiglinsky, Katherine. “Fintechs Found to Be Much More Likely to OK Suspicious PPP Loans.” Bloomberg News, August 17, 2021. <https://www.bloomberg.com/news/articles/2021-08-17/fintechs-found-to-be-much-more-likely-to-ok-suspicious-ppp-loans>.

normal detection systems but can often be caught by regtech analysis using more data and machine learning.

Many regtech solutions for fighting money laundering grew out of technology for identifying fraud, which has generally been more advanced. This may be because the industry has an enormous financial interest in preventing fraud losses. It may also reflect the fact that, in fraud, firms are usually dealing with the certainty of a problem, whereas in AML, they usually never know whether the “Suspicious Activity Reports” they file with FinCEN lead to something useful. These factors make it all the more important to equip banks and their regulators with tools that can more easily, and less expensively, detect patterns of crime.

Consumer protection and financial inclusion

U.S. consumer protection law bans Unfair and Deceptive Acts and Practices (UDAP), both in the financial sector and overall, and adds the criterion of “abusive” activity for purposes of enforcement by the Consumer Financial Protection Bureau (UDAAP). However, enforcement of subjective standards like “unfairness” and “deception” is challenging, often hampered by the difficulty of detecting and analyzing patterns of potentially illegal behavior. As with discrimination, UDAAP enforcement relies on considerable subjective judgment in distinguishing activities that are against the law from more benign patterns. This also makes compliance difficult. AI-based regtech can bring to bear the power of more data and AI analytical tools to solve these challenges, allowing regulators to detect and prove violations more easily. It might also enable them to issue more clear and concrete guidance—including more sophisticated standards on statistical modeling—to help industry avoid discrimination and being responsible for UDAAPs.

There is a growing recognition among advocates that full financial inclusion, especially for emerging markets, requires greatly expanded use of digital technology. Access to cell phones has, in effect, [put a bank branch in the hands of two-thirds of the world's adults](#). This unprecedented progress has, in turn, highlighted barriers to further success, most of which could be solved or ameliorated with better data and AI.

One is the problem of AML “de-risking.” As noted above, banks must follow Know-Your-Customer (KYC) rules before accepting new customers, a process that includes verifying the person's identity. In many developing countries, poor people—and particularly women—lack formal identity papers like birth certificates and driver's licenses, effectively excluding them from access to the formal financial system.¹³ In some parts of the world, the regulatory pressure on banks to manage risk associated with taking on new customers has resulted in whole sectors—and, in some countries, the entire population—being cut off from banking services.¹⁴ In reality, these markets include millions of consumers who would be well-suited to opening an account and do not present much risk at all. Banks and regulators struggle with how to distinguish high-risk individuals from those who are low risk. A great deal of work is underway in various countries to solve this problem more fully with AI, through the use of “digital identity” mechanisms that can authenticate a person's identity via their “digital footprints.”

A related challenge is that expanded financial inclusion has produced increased need for better consumer protection. This is especially important for people who are brought into the financial system by “inclusion” strategies and who may lack prior financial background and literacy, making them vulnerable to predatory practices, cyber scams, and other risks. Regulators are using AI chatbots equipped with NLP to intake and analyze consumer complaints at scale and to crawl the web for signs of fraudulent activity.

¹³ World Bank. “Inclusive and Trusted Digital ID Can Unlock Opportunities for the World's Most Vulnerable.” August 14, 2019. <https://www.worldbank.org/en/news/immersive-story/2019/08/14/inclusive-and-trusted-digital-id-can-unlock-opportunities-for-the-worlds-most-vulnerable>.

¹⁴ World Bank. “De-risking in the Financial Sector,” October 7, 2016. <https://www.worldbank.org/en/topic/financialsector/brief/de-risking-in-the-financial-sector>.

One example is the RegTech for Regulators Accelerator (R2A) launched in 2016 with backing from the Bill & Melinda Gates Foundation, the Omidyar Network, and USAID.¹⁵ It focuses on designing regulatory infrastructure in two countries, the Philippines and Mexico. Emphasizing the need for consumers to access services through their cell phone, the project introduced AML reporting procedures and chatbots through which consumers could report complaints about digital financial products directly to regulators.

Importantly, regtech innovation in the developing world often exceeds that in the major advanced economies. One reason is that many emerging countries never built the complex regulatory infrastructure that is commonplace today in regions like the U.S., Canada, and Europe. This creates an opportunity to start with a clean slate, using today's best technology rather than layering new requirements on top of yesterday's systems.

Credit discrimination and predatory lending

Perhaps AI's greatest financial inclusion promise lies in the emergence of data-centered credit underwriting techniques that evaluate loan applications. Traditional credit underwriting has relied heavily on a narrow set of data—especially the individual's income and credit history, as reported to the major Credit Reporting Agencies—because this information is easily available to lenders. Credit scores are accurate in predicting default risk among people with good FICO scores (and low risks of default). However, those traditional underwriting techniques skew toward excluding some people who could repay a loan but have a thin credit file (and hence a lower or no credit score) or a complicated financial situation that is harder to underwrite.

AI underwriting is beginning to be used by lenders, especially fintechs. AI is also increasingly being used by financial firms as a regtech tool to check that the main underwriting process complies with fair-lending requirements. A third process, much less developed, is the potential for the same technologies to be used by regulators to check for discrimination by lenders, including structural bias and unintentional exclusion of people who could actually repay a loan. Structural biases often lead to “disparate impact” outcomes. In these cases, regulators assert that a lending policy was discriminatory on the basis of race, gender, or other prohibited factors, not because of intent but because a specific class of consumers endured negative outcomes. Because disparate impact is a legal standard¹⁶ and violations of these laws create liability for lenders, these claims may also be made by plaintiffs representing people who argue they have been wronged.

Research conducted by [FinRegLab](#) and others is exploring the potential for AI-based underwriting to make credit decisions more inclusive with little or no loss of credit quality, and possibly even with gains in loan performance. At the same time, there is [clearly risk](#) that new technologies could exacerbate bias and unfair practices if not properly designed, which will be discussed below.

Climate change

In March 2022, the Securities and Exchange Commission proposed rules for requiring public companies to disclose risks relating to climate change.¹⁷ The effectiveness of such a mandate will inevitably be limited by the fact that climate impacts are notoriously difficult to track and measure. The only feasible way to solve this will be by gathering more information and analyzing it with AI techniques that can combine vast sets of data about carbon emissions and metrics, interrelationships between business entities, and much more.

¹⁵ Grupo Pasarela Audiovisual. Accessed May 20, 2022. <https://www.r2accelerator.org/>.

¹⁶ United States Supreme Court. *Texas Department of Housing and Community Affairs et al. v. Inclusive Communities Project, Inc.* 2014. https://www.supremecourt.gov/opinions/14pdf/13-1371_8m58.pdf.

¹⁷ Securities and Exchange Commission. “SEC Proposes Rules to Enhance and Standardize Climate-Related Disclosures for Investors.” 2022. <https://www.sec.gov/news/press-release/2022-46>.

Challenges

The potential benefits of AI are enormous, but so are the risks. If regulators mis-design their own AI tools, and/or if they allow industry to do so, these technologies will make the world worse rather than better. Some of the key challenges are:

Explainability: Regulators exist to fulfill mandates that they oversee risk and compliance in the financial sector. They cannot, will not, and should not hand their role over to machines without having certainty that the technology tools are doing it right. They will need methods either for making AIs' decisions understandable to humans or for having complete confidence in the design of tech-based systems. These systems will need to be fully auditable.

Bias: There are very good reasons to fear that machines will increase rather than decrease bias. Technology is amoral. AI "learns" without the constraints of ethical or legal considerations, unless such constraints are programmed into it with great sophistication. In 2016, Microsoft introduced an AI-driven chatbot called Tay on social media. The company withdrew the initiative in less than 24 hours because interacting with Twitter users had turned the bot into a "racist jerk." People sometimes point to the analogy of a self-driving vehicle. If its AI is designed to minimize the time elapsed to travel from point A to point B, the car or truck will go to its destination as fast as possible. However, it could also run traffic lights, travel the wrong way on one-way streets, and hit vehicles or mow down pedestrians without compunction. Therefore, it must be programmed to achieve its goal within the rules of the road.

In credit, there is a high likelihood that poorly designed AIs, with their massive search and learning power, could seize upon proxies for factors such as race and gender, even when those criteria are explicitly banned from consideration. There is also great concern that AIs will teach themselves to penalize applicants for factors that policymakers do not want considered. Some examples point to AIs calculating a loan applicant's "financial resilience" using factors that exist because the applicant was subjected to bias in other aspects of her or his life. Such treatment can compound rather than reduce bias on the basis of race, gender, and other protected factors. Policymakers will need to decide what kinds of data or analytics are off-limits.

One solution to the bias problem may be use of ["adversarial AIs."](#) With this concept, the firm or regulator would use one AI optimized for an underlying goal or function—such as combatting credit risk, fraud, or money laundering—and would use another separate AI optimized to detect bias in the decisions in the first one. Humans could resolve the conflicts and might, over time, gain the knowledge and confidence to develop a tie-breaking AI.

Data quality: As noted earlier, AI and data management are inextricably intertwined, so that acceptable AI usage will not emerge unless regulators and others solve the many related challenges regarding data use. As with any kind of decision making, AI-based choices are only as good as the information on which they rely.

Accordingly, regulators face tremendous challenges regarding how to receive and clean data. AI can deal most easily with "structured data," which arrives in organized formats and fields that the algorithm easily recognizes and puts to use. With NLP tools, AI can also make sense of "unstructured data." Being sure, however, that the AI is using accurate data and understanding it requires a great deal of work. Uses of AI in finance will require ironclad methods for ensuring that data is collected and "cleaned" properly before it undergoes algorithmic analysis. The old statistics maxim "garbage in, garbage out" becomes even more urgent when the statistical analysis will be done by machines using methods that its human minders cannot fully grasp.

It is critical that policymakers focus on what is at stake. AI that might be good at, say, recommending a movie to watch on Netflix will not suffice for deciding whether to approve someone for a mortgage or a small-business loan or let them open a bank account.

Data protection and privacy: Widespread use of AI will also necessitate deep policy work on the ethics and practicalities of using data. What kinds of information should be used and what should be off-limits? How will it be protected from security risks and government misuse? Should people have the right to force-remove past online data, and should companies' encryption techniques be impenetrable even by the government?

Privacy-enhancing technologies may be able to mitigate these risks, but the dangers will require permanent vigilance. The challenge will spike even higher with the approach of quantum computing that has the power to break the encryption techniques used to keep data safe.

Model Risk Management (MRM): Mathematical models are already widely used in financial services and financial regulation. They raise challenges that will only grow as AI becomes more widely employed. This is particularly true as AI is placed in the hands of people who do not understand how it makes decisions. Regulators and industry alike will need clear governance protocols to ensure that these AI tools are frequently retested, built on sufficiently robust and accurate data, and are kept up to date in both their data and technical foundations..

AI roadmap for regulators

Redesigning financial regulation to catch up to the acceleration of AI and other industry innovation is somewhat analogous to the shift in cameras from analog to digital at the turn of the millennium. An analog camera produces an image in a form that is cumbersome, requiring expert (and expensive) manipulation to edit photos. Improving the process of taking pictures with 35-millimeter film hits a ceiling at a certain point. By comparison, the digital or smartphone camera was a whole new paradigm, converting images into digital information that could be copied, printed, subjected to artificial intelligence for archiving and other methods, and incorporated into other media. The digital camera was not an evolution of the analog version that preceded it. It was entirely different technology.

Similarly, current regulatory technologies are built on top of an underlying system of information and processes that were all originally designed on paper. As a result, they are built around the constraining assumptions of the analog era, namely that information is scarce and expensive to obtain, and so is computing power.

To undertake a more dramatic shift to a digitally native design, regulators should create new “taxonomies” of their requirements (which some agencies are already developing) that can be mapped to AI-powered machines. They should also develop comprehensive education programs to train their personnel in technology knowledge and skills, including baseline training on core topics, of which AI is a single, integral part. Other key “big data” issues include the Internet of Things, cloud computing, open source code, blockchains and distributed ledger technology, cryptography, quantum computing, Application Program Interfaces (APIs), robotic process automation (RPA), privacy enhancing technologies (PETs), Software as a Service (SaaS), agile workflow, and human-centered design.

These are big challenges that bring substantial risks, but the cost of sticking with largely analog systems is greater. Personnel may fear that such an overhaul could result in machines taking their jobs, or that machines will make catastrophic errors, resulting in financial mishaps. On the former fear, robotics and AI can in fact empower human beings to do their jobs better, by decreasing vast amounts of routine work duties and freeing up people to use their uniquely human skills on high-value objectives. On the second fear, agencies should build cultures grounded in an understanding that humans should not cede significant decisionmaking to machines. Rather, experts should use technology to help prioritize their own efforts and enhance their work.

Data is the new oil not only in its value but in its impact: Like oil, digitization of data can solve some problems and cause others. The key to achieving optimal outcomes is to use both data and AI in thoughtful ways—carefully designing new systems to prevent harm, while seizing on AI's ability to analyze volumes of information that would overwhelm traditional methods of analysis. A digitally robust regulatory system with

AI at its core can equip regulators to solve real-world problems, while showcasing how technology can be used for good in the financial system and beyond.



The Center on Regulation and Markets at Brookings provides independent, non-partisan research on regulatory policy, applied broadly across microeconomic fields. It creates and promotes independent economic scholarship to inform regulatory policymaking, the regulatory process, and the efficient and equitable functioning of economic markets.

Questions about the research? Email communications@brookings.edu.
Be sure to include the title of this paper in your inquiry.

© 2022 The Brookings Institution | 1775 Massachusetts Ave. NW, Washington, DC 20036 | 202.797.6000