Brave new world: AI and its implications for compliance organizations

Building a better working world

It's likely that 2023 will be remembered as the moment when artificial intelligence (AI) took the leap from theoretical possibility to tangible opportunity. Just as the advent of cryptocurrency in the early 2010s ignited an explosion of interest in potential use cases for distributed ledgers and decentralized virtual assets, recent developments in deep learning and computational reasoning have sparked widespread discussion of AI's potential applications across the working world. As new AI technologies begin to enter the financial services industry, compliance professionals will be forced to rethink existing operational models and traditional approaches to risk management. They should do so with cautious optimism.

Around the corner

In November 2022, the technology company OpenAI launched ChatGPT, a chatbot using a large language AI model trained on billions of parameters, which exhibited an impressive ability to answer abstract questions, pass aptitude tests, create style-transfer text, and produce competent responses to most user prompts. Since 2019, OpenAI has partnered with Microsoft to build supercomputing platforms that enable the deep learning initiatives that produced ChatGPT; in January 2023, Microsoft doubled down on its partnership by announcing a multibillion dollar investment in OpenAI. This decision by Microsoft, one of the world's leading providers of business productivity software, suggests that generative AI models like ChatGPT will soon enter the workplace. The financial services industry — a field predicated upon speed, margin, and precision, but that still frequently operates using dated technology and labor-intensive processes — presents some of the ripest opportunities for disruption, specifically in the compliance department.

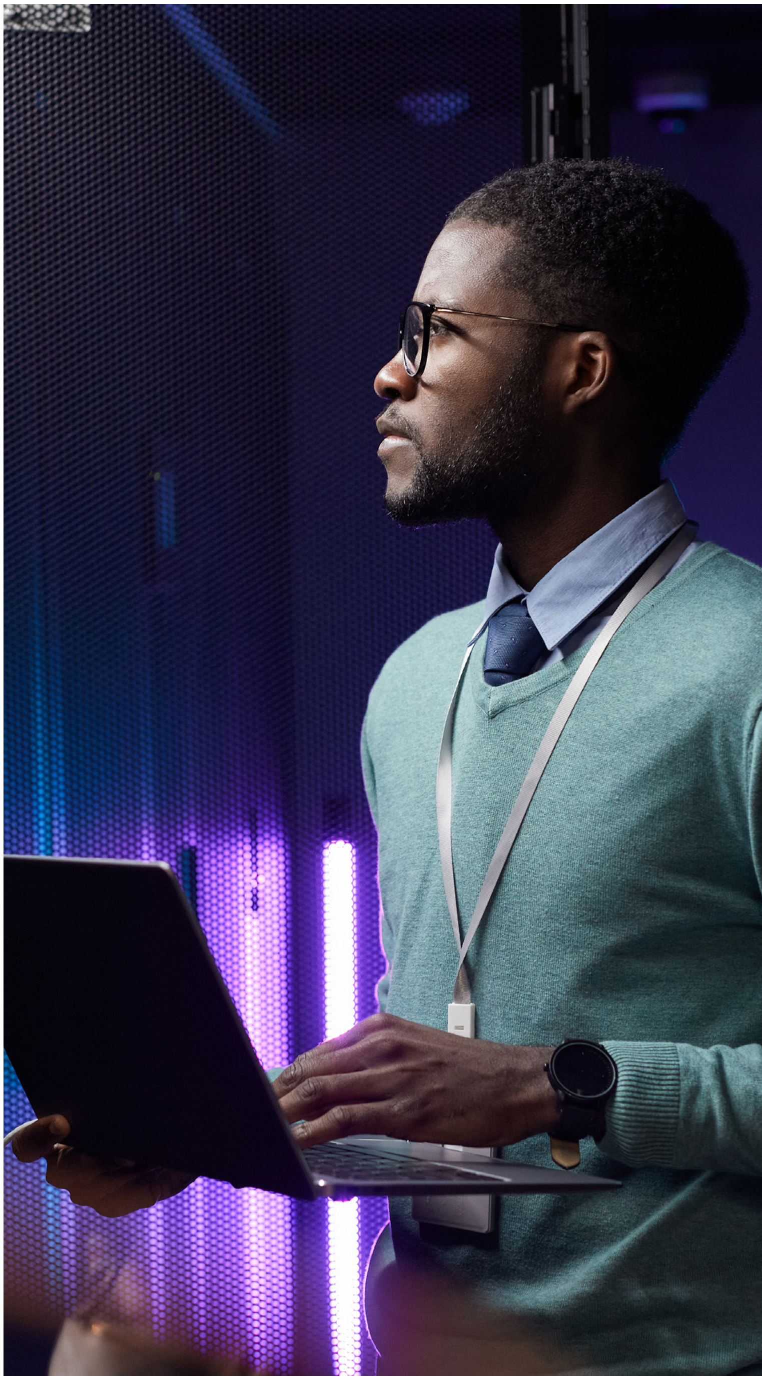

Sources:
"OpenAI forms exclusive computing partnership with Microsoft to build new Azure AI supercomputing technologies," Microsoft News Center, July 22, 2019, https://news.microsoft.com/2019/07/22/openai-forms-exclusive-computing-partnership-with-microsoft-to-build-new-azure-ai-supercomputing-technologies/; "Microsoft Invests $10 Billion in ChatGPT Maker OpenAI," Bloomberg, January 23, 2023, https://www.bloomberg.com/news/articles/2023-01-23/microsoft-makes-multibillion-dollar-investment-in-openai?leadSource=uverify%20wall.

Tangible opportunity

Historically, compliance professionals have treated technological innovation with skepticism. The financial, legal and reputational risks of punitive regulatory actions incentivize compliance departments to take a "fear-based approach" to risk management, prioritizing the avoidance of regulatory scrutiny over considerations of operational efficiency or optimization. When organizations have devoted resources to enhancing their compliance models, their focus is still on the basics: improving data quality, reducing case backlogs, managing employee turnover, etc. Deploying new technologies to streamline functions falls by the wayside, resulting in high-touch, resource- and labor-intensive processes.

In recent years, however, industry leaders have recognized the shortcomings of manual operations and begun to experiment with advanced technologies such as machine learning and intelligent automation in an attempt to improve risk outcomes while controlling costs. Common examples include automated negative news screening and analysis, automated data retrieval through robotic process automation (RPA) and application programming interfaces (APIs), automated template-based narrative generation, and leveraging predictive models to enhance or accelerate decisioning (e.g., risk scoring for transaction monitoring, and list screening alerts). As even stronger AI tools enter the industry in coming years, they will augment core components of the compliance program model, including the following:

Governance

Natural language models (like GPT-3 which undergirds ChatGPT) will be able to scan thousands of approved sources for regulatory updates and produce consolidated summaries of the most important information for review and interpretation by senior management. When regulatory changes require updates to firm standards and procedures, AI solutions will create first drafts of policy documents based on specified inputs and parameters, which can then be refined by the human eye. Indeed, a short prompt given to the current iteration of ChatGPT for experimental purposes ("Write me an AML compliance policy for a US-based financial institution") produced a remarkably accurate 10-page governance document that covered the core tenets of a compliance program. While not complete, this sample document could easily serve as the starting point for a broader procedural uplift. Relying on technology to support the governance model will enable firms to reduce the costs of key initiatives such as regulatory mapping and expedite the change management and procedural update process.

False positive dispositioning

Financial institutions currently rely on simplistic models for processes like transaction monitoring, adverse media screening, and sanctions screening to identify illicit customer activity. Without highly tuned parameters, these models produce a high volume of false positives, which must be investigated and discounted. AI tools will help resolve false positive issues in two ways: First, improved model tuning logic will better identify potential true matches. Screening for "David Johnson" might produce 1,000 low-quality hits; screening for "David Johnson" in combination with other inputs (e.g., age, address, contact information, citizenship, work history, vehicle registrations, relationship history, known associates) pulled from open and closed sources might produce, say, 11 high-quality results that better optimize the reviewer's time. Second, language models will capture and highlight the most pertinent information to an investigation, reducing overall handling time and improving material risk identification.

SAR writing

When financial institutions identify potentially suspicious activity, they are required to investigate and, if deemed material, file a suspicious activity report (SAR) with regulators. SARs generally range from one to 10 pages and are composed by investigators who combine transactional analysis and risk expertise to demonstrate potentially criminal activity to government authorities. These reports are among the most important components of a strong Bank Secrecy Act/Anti-Money Laundering (BSA/AML) compliance program, and the composition process will retain an element of human touch for the foreseeable future given the variability in risk appetite and decisioning methodology between institutions. That said, AI models today can produce standardized reports of all kinds if trained on sufficient sample data, and soon these solutions will be employed by banks to meaningfully accelerate the SAR writing process. Instead of relying on a team of investigators to draft reports or customize SAR templates with transaction analysis, AI will be used to identify the suspicious customer activity, determine the corresponding risk typology, perform open- and closed-source searches to form a coherent customer profile, and compose a report outlining customer background, suspicious activity, and rationale for filing a SAR. As opposed to manually composing a handful of reports per day, investigators will assume a quality control role enabling them to review and file SARs at scale. In turn, institutions will experience fewer alert backlogs (and resulting regulatory actions) while lowering operational costs.

Ongoing monitoring

As discussed above, financial institutions are required to monitor their customers on an ongoing basis to identify potentially fraudulent or criminal activity between normal customer review cycles. AI-based solutions will be able to revamp the detection logic undergirding screening and monitoring tools, leveraging additional data sources to create holistic customer profiles that better define suspicious activity for a specific customer and track risk across multiple domains (AML, fraud, sanctions, etc.). These shifts will help reduce false-positive alerts and better identify compliance risks, allowing institutions to migrate away from rigid periodic risk review schedules in favor of more agile ongoing due diligence models.

Regulatory perspective

Just as the benefits and pitfalls of AI to society and the economy remain to be seen, so too do the rules and frameworks that will govern its adoption. The recent case study of digital asset regulation indicates that regulators often take a wait-and-see approach to nascent technology, with guidance trailing innovation by three to five years. While it's impossible to predict the shape of the regulatory overlay, there are a few main themes inherent to AI with which financial regulators will have to grapple.

AI bias

AI models are only as good as the data they are trained on, including human feedback provided to improve their performance. Natural language models in particular often use reinforcement learning, a technique in which humans serve as "labelers" to validate model outputs and identify correct answers. Without guardrails, this human feedback can, however unintentionally, introduce bias to the model's data, with downstream impacts to AI decisioning. Regulators in the future will need to ensure that the reliance on AI for business support does not result in unequal distribution of its benefits.

Explainability

Financial institutions operating within a regulatory environment are often called upon by regulatory authorities to substantiate their risk decisions. Decision-making enabled by complex AI tools governed by thousands of underlying indicators may help to accelerate certain processes, but institutions must be mindful of their (or their vendors') ability to produce rationales for these decisions in a format that can be understood/interpreted by nontechnical resources, such as relevant citations for a fact-based search or key data attributes/values influencing predictive model outputs. Regulators will look to establish minimum standards for risk decisioning, and it will be incumbent on compliance personnel to understand requirements and maintain sufficient line of sight into the information sources and logic utilized by AI models to offer recommendations or support decision-making.

Data management

One immediate advantage of leveraging AI tools will be enhanced customer insights, combining inputs from open and closed sources to create advanced profiles of risk, behavior and profitability. These profiles, however, are predicated upon the availability, quality and security of customer data. In an AI economy, personal data comes at a premium, becoming even more valuable than it is today. In response to commercialization of personal data by large technology firms in recent years, regulatory bodies around the world have already moved to implement data privacy and security laws designed to return control to the individual the data pertains to. Efforts to secure, protect and govern access to data will only accelerate as AI is deployed commercially and demand for personal data grows.

Cyber risks

As the economy becomes ever more dependent on technology and data, the potential for hacks and data breaches will increase. Furthermore, as advanced technologies become more accessible to the general public, use of these technologies to execute scams and other fraudulent activities is likely to become more common (e.g., using AI to alter pictures or create fake videos to influence the behavior of unsuspecting consumers). Financial institutions today struggle to protect against cybercrime, and regulators have implemented cybersecurity laws to govern the strength and durability of these controls. As AI is introduced into the economy at scale, the potential for illicit actors to access and manipulate AI models, along with their underlying data, will become an even greater concern. Regulators will be tasked with revamping cybersecurity frameworks to account for these incremental risks and combat the unauthorized use of AI tools for personal gain; compliance organizations will need to monitor and adapt accordingly.

View from the corner office

Despite their promise, AI tools are in the early stages of the technology adoption curve. Training remains costly and demands a large number of data points to enable continuous learning; model outputs may require validation or third-party verification. AI solutions have not yet demonstrated a level of dependability necessary to permanently enhance human decision-making, much less displace it. Moreover, the indeterminant nature of compliance risk management — how often does a SAR truly prevent criminal activity? — impedes the ability to refine model outcomes via reinforcement learning methodologies.

Over the next decade, technologists will need to coordinate with public and private sector stakeholders to remove bias, reduce costs and improve reliability. In turn, compliance professionals should approach AI with hopeful skepticism and take a few key actions:

1. Challenge existing models.

As AI tools start to become viable from a cost and value perspective, compliance professionals should perform top-down and bottom-up assessments of their operating model to proactively identify areas for potential enhancement. Designing a future-state compliance framework with key operational objectives — risk outcomes, cost reduction, process improvement — can help govern the implementation journey.

2. Prioritize customer data.

In a process and control environment operated by AI, customer data is paramount. Although standardization of AI solutions is a few years away, compliance leaders should begin enhancing their customer data now. Data remediations are complex and lengthy, frequently spanning multiple years. Taking steps in the near future to improve the availability and quality of customer data down the line will pay dividends, not only in the feasibility of long-term technology integrations, but in the short term through improved customer insights and risk management.

3. Focus on process augmentation.

The concept of AI in the workplace inevitably inspires concerns over staffing changes. While a day may come when human intelligence becomes obsolete, we don't have to worry about crossing that bridge for quite some time. In the near to medium term, AI tools will only succeed in augmenting certain processes, almost all of which will still require review and/or validation by a human. Even the most advanced AI models — the tools that can investigate, analyze, report, compose and manipulate data — can frequently be wrong. The most popular and innovative AI training methodology currently in use, deep-learning, quickly struggles when faced with an input outside its original parameters. The result is a very powerful tool that can be relied upon to automate certain components of a process (e.g., drafting documentation, scouring the internet to disprove false-positive alerts), but cannot yet be trusted to perform that process from end to end (e.g., compliance policy design, customer risk decisioning). While automation may displace some roles, its efficiencies at scale will create others. In the coming years, compliance professionals should prioritize ways in which AI can improve, rather than supplant, processes and controls.

4. Pursue relevant skill sets.

AI-enabled automation will create a demand for people who understand the logic utilized by AI tools and can craft effective inputs to yield the most meaningful responses. Similar to the proliferation of "search engine optimization" jobs in the early 2000s, the deployment of AI in compliance organizations will necessitate a new type of technological fluency that is able to maximize the utility and quality of AI model outputs — including the ability to decipher when an outcome may be incorrect. Compliance leaders should ensure they are fostering the relevant skill sets and attracting the requisite talent to optimize the efficacy of AI tools.

Conclusion

As the world reacts over the latest advancements in AI, regulators and business leaders must consider the practical opportunities and risks of its implementation. When coupled with a robust regulatory regime and thoughtful application, AI technology has the potential to revolutionize almost all areas of the financial services industry. Those in the compliance seat should embrace its potential and remain creative as they guide its adoption.

EY | Building a better working world

EY exists to build a better working world, helping to create long-term value for clients, people and society and build trust in the capital markets.

Enabled by data and technology, diverse EY teams in over 150 countries provide trust through assurance and help clients grow, transform and operate.

Working across assurance, consulting, law, strategy, tax and transactions, EY teams ask better questions to find new answers for the complex issues facing our world today.

EY refers to the global organization, and may refer to one or more, of the member firms of Ernst & Young Global Limited, each of which is a separate legal entity. Ernst & Young Global Limited, a UK company limited by guarantee, does not provide services to clients. Information about how EY collects and uses personal data and a description of the rights individuals have under data protection legislation are available via ey.com/privacy. EY member firms do not practice law where prohibited by local laws. For more information about our organization, please visit ey.com.

Ernst & Young LLP is a client-serving member firm of Ernst & Young Global Limited operating in the US.

US SCORE no. 19733-231US
2023-4245267
ED None

This material has been prepared for general informational purposes only and is not intended to be relied upon as accounting, tax, legal or other professional advice. Please refer to your advisors for specific advice.

EY member firms do not practice law where not permitted by local law and regulations. Ernst & Young LLP (US) does not practice law or offer legal advice.

ey.com