MAYER | BROWN

Document

Artificial Intelligence & Financial Services

Thought Leadership

Contents

How to Invest in AI Fintech: Factors to Consider..... 3

The American Data Privacy and Protection Act: Is Federal Regulation of AI Finally on the Horizon?..... 15

Explainability Is an AI Tool's Best Defense..... 20

US Insurance Regulators' Evolving Perspectives on Artificial Intelligence..... 28

UK Government proposes a new approach to regulating artificial intelligence (AI)..... 33

How to Invest in AI Fintech: Factors to Consider

Jennifer Carlson, Joe Castelluccio, Nina Flax and Elizabeth Raymond

Summary of Key Points

- Artificial intelligence ("AI") is one of several technologies being developed and deployed across fintech sectors.
- Transactions that enable these activities can take a variety of forms, including joint ventures and strategic partnerships, minority and majority investments, M&A-style acquisitions and public listings.
- Financial institution investors should first define their AI goals and strategy, and then attempt to align their investment tactics with their AI strategy. As these AI strategies evolve, so will the transactions for investing in AI.
- Whether you are a financial institution, emerging fintech company or an investor focused on fintech, this article provides an assessment of the structural, risk and legal considerations to balance in these transactions.

Artificial intelligence has long since graduated from science fiction and speculative use cases to widespread and diverse development across fintech sectors. *CB Insights* reported that financings for AI startups showed continued robust activity through Q2 2022, despite changes in market sentiment from the exuberance of 2021.[1] Financial institutions of all types, and their business partners, customers and investors, view AI as an essential element of their long-term goals of staying competitive and relevant in a rapidly changing market.

Transactions among market participants can provide the missing piece each party needs to catapult its strategy forward. Different market participants bring different advantages and needs to the table, which impact the types of transactions each may employ at any given time. For example, large financial institutions often have financial resources, expertise to manufacture compliant financial products and a wealth of data about their customers' financial activities. On the other hand, companies seeking to disrupt in this space, including those offering AI products and services relating to financial services ("AI fintech companies"),bring a singular focus on a sector or product and the flexibility of a small and nimble organization but often need resources to achieve scale and viability for their products. At the same time, financial sponsors and other investors in the space often have a higher risk appetite and the ability to deploy capital quickly. However, these investors also need to monetize their investments periodically.

While there is no one-size-fits-all approach for developing and deploying AI in new/updated financial products and services, there is a common set of tools that can be used to collaborate, invest in and monetize AI in business models. Given the increasing speed at which AI and fintech are developing, the older sourcing strategies of "build versus buy" are being replaced with strategies that allow for flexible and rapid collaboration across a variety of acquisition models.

1 CB Insights, State of AI – Global Q2 2022, pages 8 and 17

This article will examine some of the strategies or "tools" on the spectrum of AI investment and monetization events: joint ventures and strategic partnerships, minority and majority investments, M&A-style acquisitions and, finally, public listings. An illustration of the full spectrum of transactions and arrangements is shown below:

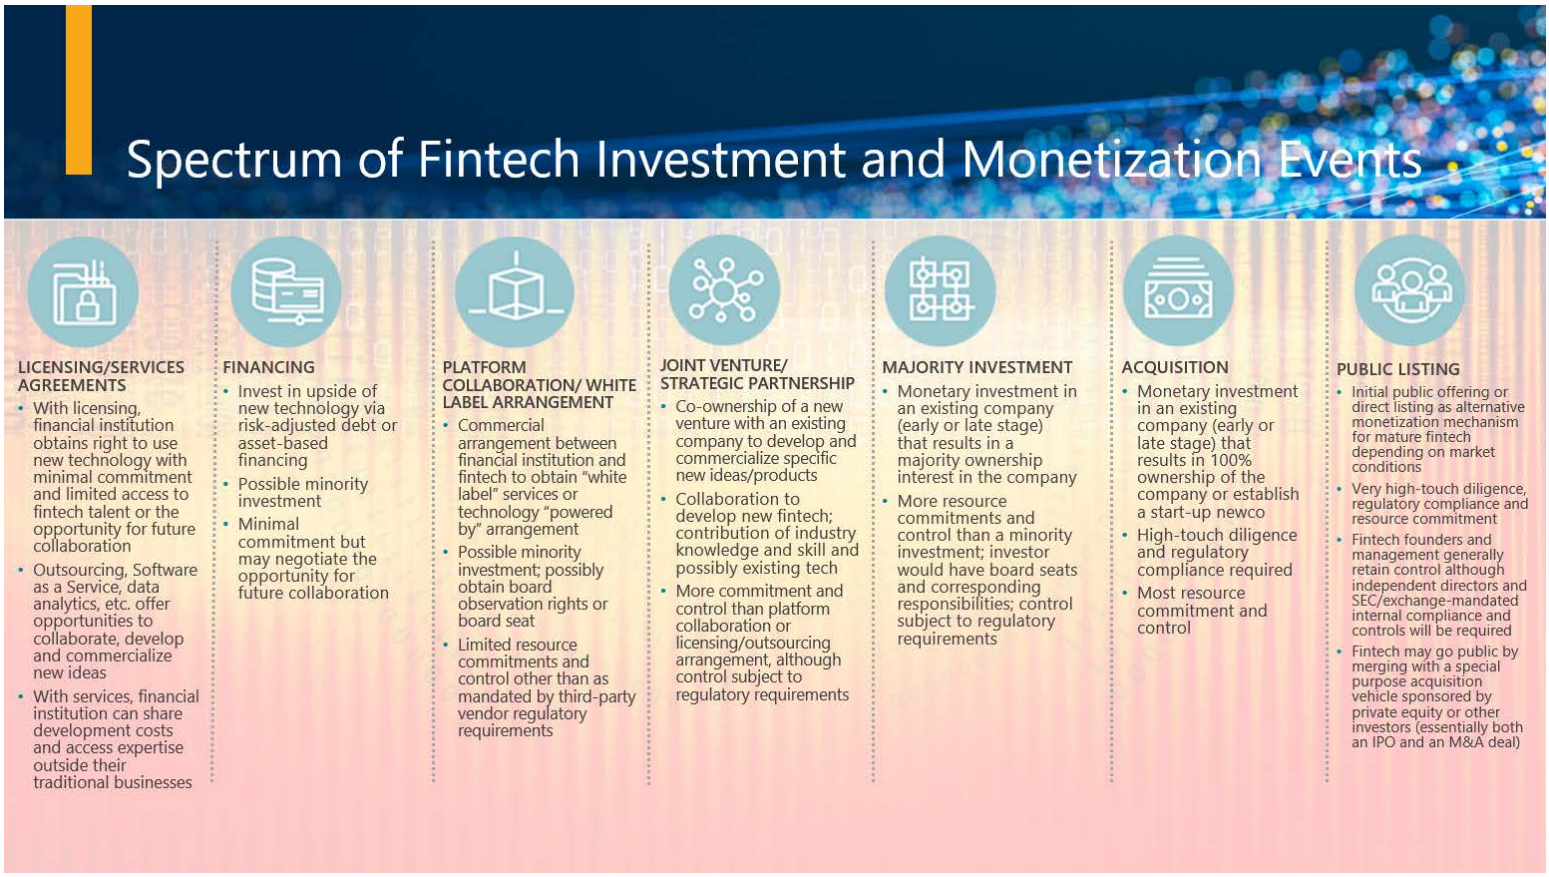

We will outline some of the due diligence, structure and contractual considerations for each type of transaction. We will focus on these considerations from the point of view of the buyer of, or the investor in, an AI fintech company. We will also review a public listing as an alternative monetization opportunity, whose availability and attractiveness will depend on market conditions.

Joint Ventures and Strategic Partnerships

The term "joint venture" is quite broad and is often used colloquially. Depending on the goals of the venture and its owners, it can involve creating a new entity, an ongoing contractual relationship or a combination of both. As distinguished from a strategic investment or an M&A transaction, a joint venture typically involves two or more parties that come together to achieve a common goal for profit.

Regulatory hurdles make it less likely for large financial institutions to joint venture or partner with an AI fintech company in the traditional sense, but other large, non-bank finance companies may consider the joint venture structure more attractive and feasible. As discussed below, a large financial institution, such as a bank holding company or an insurance company, is typically highly regulated and seeks to avoid obtaining "control" of the AI fintech company, in most cases by keeping a minority equity investment below 5 percent (or 10 percent in the case of an insurance company) of the AI fintech company's voting shares and otherwise avoiding indicia of control. Indicia of control include holding a voting seat on the company's board of directors, having certain veto or consent rights, entering into a management agreement or entering into significant business or commercial relationships with the AI fintech company. If the financial institution seeks a control relationship, it may be simpler to acquire complete control through an acquisition as opposed to a joint venture or partnership. On the other hand, the financial institution may forego any equity investment in order to avoid these control questions and seek only a commercial or financing arrangement as discussed elsewhere in our AI & Financial Services Symposium materials.

Advantages of Joint Venture. Assuming that the joint venture partners are willing to have their joint venture entity be treated as a regulated entity or the joint venture entity is otherwise not subject to what may be viewed as burdensome bank or insurance regulations, there can be a number of advantages to using a joint venture entity as opposed to a contractual joint venture. These advantages include:

a) The joint venture will have access to technology, subject matter experts such as data scientists, and products contributed to the joint venture as well as distribution channels and markets with greater economies of scale.

b) A joint venture allows regulatory risks that accompany financial institutions to be shared by the joint venture partners, especially when entering a new market.

c) Internal and external constituencies (e.g., employee talent in the joint venture and end users of the technology) will perceive a separately identifiable and visible enterprise conducting the joint venture business, with the venture lending itself more to AI innovation than to regulated bank or insurance activity.

d) Interests in a joint venture are generally easier to sell or transfer than a collection of contractual relationships.

e) The joint venture entity creates an independent vehicle with greater flexibility and convenience for capital-raising activities.

f) The joint venture entity provides a familiar structure (e.g., a corporation, limited liability company or limited partnership) in which management and governance rules can be established and in which directors, officers and employees typically play familiar roles in making decisions and implementing them, with this level of oversight likely being important in the developing area of AI.

g) The joint venture entity provides a convenient vehicle for measuring profits and allocating and distributing them to the joint venture parties.

h) The joint venture entity can be an independent employer providing identification and focus for employees, including incentive

compensation such as equity interests and the opportunity to work on cutting-edge AI projects.

i) The joint venture entity largely enables the joint venture parents to insulate themselves from the liabilities of the joint venture business.

j) The joint venture entity creates the potential for flexibility in addressing tax matters.

Disadvantages of Joint Venture.
Disadvantages of the joint venture structure, in addition to the perhaps overriding concern that the AI fintech joint venture will become a regulated entity based on its control by a regulated financial institution, include:

a) The joint venture structure will likely be more complex because establishing a separate joint venture entity often involves initial and ongoing issues, tasks and costs that are not necessarily present in a contractual association, with time-consuming oversight required by senior managers of the alliance participants.

b) The unwind process will likely be more complicated because assets, contracts, employees and other resources of the joint venture business may be property of, or affiliated with, the joint venture entity.

c) The joint venture partners may lose control because the joint venture business will normally be, in large part, conducted by the joint venture entity and the rights and ability of the joint venture entity and its activities will be limited by the governance rules of the joint venture entity.

d) Difficult fiduciary duty and conflict of interest issues may arise with a joint venture entity that may not arise in a contractual joint venture (although these can largely be handled contractually).

e) The contractual joint venture can allow more flexibility in staging and developing the joint venture by establishing an initial "let's get our feet wet" relationship without the more substantial commitment involved in establishing, and providing assets and other resources to, a separate joint venture.

Minority Investments and M&A Transactions

Strategic investments and M&A transactions offer a large financial institution, such as a bank or insurance company, some additional flexibility to tailor an investment to its specific business strategy, with each structure having its own unique advantages and disadvantages. Two general concerns applicable to each structure are (1) the "control" analysis described above in the "Joint Ventures and Strategic Partnerships" section and the effect of bank or insurance regulatory control on the AI fintech company and (2) the level of diligence a potential investor should complete with respect to each structure. In this section, "investor" refers to financial institutions as investors in or acquirers of AI fintech companies.

Advantages of Minority Investment. A passive, non-controlling investment can offer a large financial institution investor and the AI fintech company a number of advantages. These advantages include:

a) allowing the investor to leverage the AI offerings of the AI fintech company in its business with relatively low risk to the investor due to a limited commitment of resources;

b) potentially less stringent due diligence requirements of the AI fintech company, in

general, than majority investments and M&A transactions, but this can vary depending on the cost/benefit analysis and risk tolerance of each individual investor;

c) the imposition of fewer regulatory burdens on the AI fintech company;

d) allowing the AI fintech company to leverage the infrastructure and expertise of the investor; and

e) the AI fintech company's retention of a certain level of autonomy.

Disadvantages of Minority Investment. Disadvantages of this structure include:

a) very limited investor control over the AI fintech company's activities (e.g., generally no board seat or board observer, few consent rights over activities of the AI fintech company, etc.);

b) limited investor protective provisions (e.g., lead investor controls vote, future capital raises possibly diluting and subordinating the investment, etc.);

c) requiring the investor to conduct a relatively complex and ongoing control analysis for regulatory purposes; and

d) tension created due to the differing goals of the investor (financial return) and the AI fintech company (long-term viability).

The obligation of the investor to continually assess its level of control over the AI fintech company to avoid subjecting the AI fintech company to regulatory oversight is a key disadvantage to a minority investment. For example, a bank holding company investor must ensure its equity investment remains below 5 percent in addition to monitoring other means of exercising control over the AI fintech company, such as the appointment of a board member, veto rights over certain actions of the AI fintech company, ownership of 25 percent or more of any class of voting securities, rights of first refusal and ownership of convertible securities.[1] As a protective measure, a minority bank holding company investor should seek to include certain transfer rights, such as a put right, for itself in connection with its investment to allow the investor to exit the AI fintech company if regulatory concerns arise.

Investments by insurance companies (or their affiliates) will potentially be subject to the laws governing insurance holding company systems in the states where the insurance companies are domiciled (or deemed commercially domiciled). Generally, those laws presume control, and thus an affiliate relationship, to exist where one person, directly or indirectly, owns 10 percent or more of the voting securities of another person, although that presumption can be rebutted by submitting a disclaimer of control to the domiciliary state insurance commissioner. In addition, other types of rights, such as the appointment of board members, may be deemed by an insurance commissioner to constitute control of an entity, so being below 10 percent should not be considered a "safe harbor" that automatically negates control. The laws in many states limit the ability of an insurance company to acquire a controlling minority interest in another entity. In addition, if an entity

[1] Note that a potential alternative path for a bank holding company that has elected "financial holding company" status to invest in AI fintech companies is under the merchant banking authority in section 4(k)(4)(H) of the Bank Holding Company Act. This article will not attempt to address merchant banking authority, in part because its requirements (including with respect to the "routine management or operation" of a merchant banking portfolio company) are relatively restrictive.

is treated, for insurance regulatory purposes, as an affiliate of an insurance company, that relationship will need to be disclosed in the insurance company's statutory financial statements and annual holding company registration statements and enterprise risk reports, and the domiciliary state insurance commissioner will need to be notified in advance of material transactions between the insurance company and its affiliate, giving the commissioner an opportunity to review the transaction before it can go into effect.

Factors to consider when making a non-controlling, minority investment include: (a) when the financing round occurs in the lifecycle of the company (i.e., an early stage financing such as Series A or a later stage financing round such as Series D); and (b) whether the investor is the lead investor in the financing round. With respect to the financing round, investors should consider the potential regulatory risks that the AI fintech company may face because many startups operate with a focus on growth rather than regulatory compliance, and without a well-developed regulatory framework, it may be harder to ascertain the regulatory risks in an early-stage company as opposed to a more mature AI fintech startup. Further, lack of liquidity and valuation risks are common with all startups regardless of industry because these companies are hard to value when they are immature. On the other hand, investing into a later stage financing round also comes with its own issues and considerations. While more mature companies have potentially lower compliance risk exposures and more well-developed valuation methodologies, investing into a later stage financing round means entering into a more crowded capitalization table with lower returns upon an exit and potentially less control and protections than earlier investors have.

When making a minority investment, it is important to consider the benefits and disadvantages of being the lead investor in a financing round. A lead investor is generally the investor that is making the largest investment in the round, and therefore is the investor that is in charge of negotiating the terms of the round, coordinating the process between the parties, and conducting the most extensive diligence. The benefits of being a lead investor generally include the creation of a strong relationship with the AI fintech company due to the amount of interaction during, and after, the financing transaction. Additionally, the lead investor typically will have the right to negotiate a side letter with the AI fintech company, which will include contractual terms separate from the terms generally agreed upon in the financing related to the equity being purchased, and may have the advantage over other investors for negotiating a strategic alliance.

Advantages of Majority Investment/M&A Transaction. Considerations to keep in mind when determining whether to be a lead investor include:

a) the size of the investment in order to be considered the lead investor,

b) the expenses associated with conducting the diligence process and engaging various advisors as part of the process and

c) whether there are any other commercial relationships between the parties. The size of the investment can become substantial, especially in later stage financing rounds, and when there are potential venture capital funds involved, where a significant investment may

be required in order to obtain the related benefits.

Similarly, the more complex the business, the more costly the diligence and negotiation processes become. A passive non-lead investor may be able to obtain most of the same benefits without incurring as many costs. Finally, if there are commercial relationships between the AI fintech company and an investor or if the investor must abide by specific regulatory restrictions on its investment, the investor may be able to negotiate additional contractual rights pursuant to a side letter without having to be the lead investor of the financing round. For example, it is typical for a passive bank holding company investor to put its bank regulatory representations and covenants in a side letter.

Alternatively, if a large financial institution seeks a control relationship, it can structure its investment as a majority investment or an M&A transaction. Some advantages of a majority investment include:

a) providing more investor control over the AI fintech company than in a minority investment;

b) allowing the investor the opportunity to enhance the operational efficiency of the AI fintech company and address any existing risks (e.g., amend existing material agreements to address deficiencies); and

c) providing the AI fintech company with a greater opportunity to leverage the infrastructure and expertise of the investor.

Disadvantages of Majority Investment/M&A Transaction. Disadvantages of a majority investment include:

a) subjecting the AI fintech company to regulatory oversight;

b) requiring a much larger resource commitment from the investor, which entails a higher level of risk, necessitating a much higher level of due diligence (raising the issue of whether it may be more advantageous to acquire the entire AI fintech company);

c) requiring a higher level of investor responsibility and oversight with respect to the operations of the AI fintech company, including regulatory compliance; and

d) integration issues with respect to the cultures of the investor and AI fintech company. The effect of the investor obtaining control of the AI fintech company is one of the most important factors for the investor's consideration.

Generally, majority investments require a much more thorough due diligence investigation of the company than minority investments. The investor will need to assess the AI fintech company's current operations and marketing strategies (including the AI fintech company's website) and review its contracts, in each case with a particular focus on data security and regulatory compliance, as discussed more fully below. In extreme cases, it may be necessary to shut the AI fintech company down for a period of time to resolve any major issues identified in due diligence.

Lastly, a large financial institution may wish to acquire full ownership of an AI fintech company in an M&A transaction. Each of the advantages and disadvantages of a majority acquisition apply to an M&A transaction, often to a greater extent. A key additional advantage of an M&A transaction is the flexibility provided, more specifically the opportunity to use a number of different structures to address specific risks (e.g., the use of an asset sale to protect against pre-closing liabilities). Some key disadvantages of

M&A transactions include: (a) an M&A transaction requires the highest level of due diligence; and (b) concerns related to retention of key employees are at their peak.

The buyer's due diligence of an AI fintech company in an M&A transaction should include a confirmation of ownership of intellectual property and software, a personnel assessment, and an evaluation of regulatory, cybersecurity and data privacy risks. Analyzing the source code underlying the IP is critical. Open source code licenses may require disclosure to the public domain of all or a portion of the source code into which the open source code subject to any such license was incorporated. To reduce its risk, the M&A buyer should also seek to negotiate strong seller representations in the transaction documents with respect to matters such as ownership of IP, outbound licenses of the IP, use of open source code, the formatting of the source code (i.e., that it has been documented in a manner that enables a programmer of reasonable competence to understand it, manipulate it, etc.), compliance with cybersecurity and data protection laws and best practices, and other similar matters.

Due diligence of the technology and software of an AI fintech company may present amplified confidentiality issues given the competitively sensitive nature of the information being evaluated. A seller may not regard a traditional confidentiality agreement as sufficient to protect its interests where software is a large portion of the value of the business being sold. A "clean room" confidentiality agreement may be appropriate where only certain "clean team" members are permitted access to sensitive data. This is a familiar technique borrowed from transactions where antitrust issues exist because the buyer and seller are competitors. A key issue will be which individuals are permitted on the clean team. For example, will any of the buyer's employees be given access or will the team only consist of individuals employed by legal advisors and technology consultants? Will the buyer be permitted to review reports prepared by these advisors and consultants or will additional restrictions on use or redaction be required?

The buyer of an AI fintech company should also seek to address due diligence issues and risks that are particular to AI providers through targeted representations and covenants. For example, the buyer should include compliance with law representations and covenants that allocate strict liability to the seller for machine learning output regardless of whether any breach is "intentional" or "negligent" or is known by the seller. Particularly where the AI fintech company engages in lending or making underwriting decisions, the buyer should address liability for discrimination and fair lending compliance, including for any disparate impact. The buyer may also seek a representation that decisioning criteria are "explainable" or at least diligence the design criteria of the AI fintech company for explainability. Cybersecurity and data privacy representations and covenants may also need to be augmented in light of data-intensive AI systems.

As part of its due diligence process, the M&A buyer should identify key employees to retain following the closing. As mentioned above, there may be substantial differences between the cultures of the financial institution buyer and the AI fintech company. Employees will often be moving from a relatively autonomous position with modernized infrastructure at the AI fintech company to a much more structured environment, often with restrictive and outdated legacy infrastructure, at the buyer. Considering

the importance of key employees, such as lead software engineers, to the AI fintech company, the buyer should ensure it is offering attractive compensation packages to encourage these employees to remain following the closing.

The buyer may seek to impose covenants in an M&A transaction that obligate the AI fintech company to address certain issues prior to closing, such as requiring the AI fintech company to bring its operations into compliance with data protection laws (including implementing any necessary changes to its IT systems), engaging a consultant to undertake a review of open source code, making changes to its marketing materials, obtaining any additional state or third-party licenses to operate the business, or renegotiating or terminating certain problematic contracts. For example, the buyer may seek to engage a consultant to undertake an information security due diligence assessment between signing and closing of the transaction. The seller may require the consultant to agree to access and use protections directly with the seller even though the assessment will be performed for the buyer. Similar concerns as discussed above in connection with a "clean team" confidentiality agreement will arise. If deficiencies are discovered during the assessment, the buyer may seek a remediation plan with specific remediation steps required prior to closing. If only a portion of the seller's business is being sold and some employees (such as founders or programmers) are being left behind, the buyer may seek strong non-compete protections so that the intellectual property and software it is buying cannot be replicated by the seller after closing.

Depending on the M&A buyer's leverage, it should also consider including closing conditions related to technology matters to avoid being forced to close the acquisition and make these changes itself post-closing, which shifts the risks associated with any necessary shutdown to the buyer. If the seller is not selling its entire business, the buyer may require the seller to separate and stand up the technology of the purchased business prior to closing. The buyer may also seek to escrow the key software code it is purchasing, and any updates or enhancements to the code, in order to ensure that the source code remains intact exactly as reviewed during due diligence and with only those changes agreed by buyer and seller.

Public Listings

Depending on market conditions, the AI fintech company may eventually decide to go public through an initial public offering ("IPO") instead of selling to an M&A buyer. The AI fintech company could also achieve public listing status through a direct listing ("DL") or a transaction with a special purpose acquisition company ("SPAC"). Public listing, whether through an IPO, DL or SPAC transaction, is viewed by earlier stage investors in companies, including in the fintech space, as an attractive exit strategy when market conditions are good and allows the AI fintech company and its investors the benefit of a larger pool of money available in the public capital markets. An IPO can also be attractive because it does not require the founder or majority owner of the AI fintech company to give up complete control.

A public listing can help both with monetization of an investor's equity interest as well as capital raising for the AI fintech company in connection with and after the listing event. One reason an AI fintech company may opt to go the public listing route instead of doing an M&A exit is because the public markets may offer a higher earnings

multiple and enterprise value than an M&A buyer is willing to pay. The drawback of a public listing for existing stockholders is that it may take time to sell their shares in full, whereas existing stockholders can fully exit at the time of the consummation of the M&A transaction.

Before undertaking a public listing, an AI fintech company should ensure it has (i) a control structure in place and a leadership team with public company experience, (ii) the capabilities and organizational infrastructure to practice financial discipline and comply with public reporting requirements, (iii) processes and personnel to track compliance with the current and evolving regulatory landscape, and (iv) the capability to manage and mitigate operational, financial and regulatory risks. These four items will be scrutinized by potential investors in a newly public AI fintech company and by underwriters during the process of preparing for the IPO, including through due diligence. These factors may be particularly important for AI fintech companies due to the nature of fintech businesses. Fintech companies tend to be more highly regulated than companies in other industries due to their handling of money and data.

Because of their role in advising on and facilitating financial transactions and creating related products, AI fintech companies may have more complicated governance, disclosure and internal controls. Investors, regulators and auditors will expect AI fintech companies to have appropriate management and systems in place to monitor operational effectiveness, allow for reliable and timely financial reporting and other public disclosure, and provide for consistent regulatory compliance. In order to facilitate this, an AI fintech company will need to have documented procedures related to corporate governance, disclosure controls and internal control over financial reports that clearly outline roles and responsibilities for company leadership and employees.

The AI fintech company should ensure that it has leaders at the board of directors and officer level with public company experience, particularly in the financial services sector, who will have the familiarity and experience to guide the newly public company on business strategy and regulatory compliance. The company will also need a finance team in place that can achieve and maintain the profitability and timely disclosures expected of a public company. Ideally the AI fintech company's compliance function will be sufficiently built out so as to allow it to identify, assess, test and monitor new regulations from US regulatory bodies such as the Federal Reserve, FDIC, OCC (or other applicable bank regulators), Consumer Financial Protection Bureau, applicable state regulators, and Securities and Exchange Commission (SEC) as they are proposed and adopted. Having these structures and personnel in place will provide comfort to potential investors and allow them to focus on the AI fintech company's investment thesis and technology instead of on concerns over regulation and compliance.

Prior to a public listing, the AI fintech company, like other companies planning to go public, will need to establish a compelling investment thesis and demonstrate that it has foreseeable revenue growth and profitability and processes for making quarterly forecasts and reporting financial results. For a public listing with concurrent financing, the AI fintech company will select underwriters or placement agents that typically are banks with which the company has a pre-existing relationship. One gating item may be that the AI fintech company's audits need to

be conducted in accordance with Public Company Accounting Oversight Board standards.

Once the public listing process kicks off, the first major undertaking is due diligence. Due diligence will help the AI fintech company, the banks (if any) and their respective counsel prepare the registration statement that registers the offering and sale of securities for the IPO, DL or SPAC transaction. The due diligence process will also help identify any issues that the AI fintech company may need to address before it can become a public company. Common areas where issues sometimes arise include ownership structures, shareholder agreements and compensation arrangements. Ultimately, the AI fintech company's auditors and officers will need to make certain assurances about the accuracy of financial information contained in the registration statement, and the process of verifying such information is completed through the due diligence process. Diligence will include providing documentation to verify data appearing in the registration statement as well as interviewing a company's management and auditors and third parties such as customers and suppliers.

The key components to the registration statement, which will be informed by the diligence discussed above, include the (i) management's discussion and analysis (MD&A), (ii) business overview, (iii) risk factors and (iv) description of offered securities. The MD&A will discuss the AI fintech company's financial results and condition. The business overview will describe the AI fintech company's business in detail, including information about the company's key products and services, competitive strengths, properties and human capital. In the risk factors section, the AI fintech company will describe the risks and challenges that may impact its financial condition and operations, including operational risks, regulatory and compliance risks, industry risks and risks related to the transition from a private to public company. For most IPOs and DLs, companies have the option to confidentially submit the registration statement to the SEC to obtain feedback prior to the registration statement being made publicly available.

Once diligence is complete and the registration statement is cleared by the SEC, the last step for a public listing depends on whether the mechanism is an IPO, a DL or a SPAC transaction.

In an IPO, the underwriters will market the offering in a process called a road show. This process will involve meetings between management, usually a company's CEO and CFO, and potential institutional investors. The company also makes a recorded version of the road show publicly available for potential retail investors. The road show typically takes about two weeks. Through this process, the underwriters track indications of interest from potential investors to help gauge the demand for the stock being sold in the offering. The underwriters will then make a pricing recommendation (comprised of how many shares can be sold and at what price) to the AI fintech company. If the board of directors approves the pricing recommendation, the underwriters will purchase all offered shares at a discount from the company and immediately resell the shares to investors at the agreed upon price to the public. Unless the IPO includes a secondary component, existing investors continue to hold "restricted securities" and usually agree to a lock-up period of 180 days. Once the lock-up expires, existing investors usually may sell shares into the market pursuant

to the exemption from registration contained in Rule 144.

In a DL, the AI fintech company's stock is listed on a stock exchange without the participation of underwriters or placement agents because there is typically no concurrent financing. A DL may be desirable for a company that does not require additional capital but that wants to provide liquidity to a large investor base. Therefore, the registration statement for the DL is a resale registration statement for existing investors. In some circumstances, the AI fintech company will hold an "investor day" prior to the DL that is similar to a roadshow in order to provide potential public investors with information about the company that is consistent with the registration statement, including financial guidance. Once the SEC declares the resale registration statement effective and the stock exchange approves the listing, existing investors may sell stock directly into the market.

In a SPAC transaction, the AI fintech company merges with a SPAC that has already completed an IPO and is an SEC reporting company. The AI fintech company and the SPAC will enter into a business combination transaction that is publicly announced and after that file a registration statement to register the offer and sale of the SPAC's securities in exchange for the AI fintech company's securities.[2] If the AI fintech company needs to raise capital, including to cover potential redemptions of SPAC shares, the companies may engage a bank to act as placement agent for a private investment in public equity (PIPE) transaction. The PIPE shares will typically be registered for resale shortly after the SPAC transaction is completed. Once the registration statement for the business combination transaction is finalized, the stockholders of both companies will need to approve the transaction. Then the companies will consummate the merger and the AI fintech company's stockholders will receive freely tradable shares, subject to any agreed lock-up period.

In each case (IPO, DL or SPAC transaction), the AI fintech company's stock will commence trading, and the company will begin its next chapter as a public company.

Conclusion

As shown in our discussion above, transactions involving investments in AI fintech companies include a wide spectrum of possible structures, with legal and business issues that vary based on the transaction type. Financial institution investors should first define their AI goals and strategy and then attempt to align their investment tactics with their AI strategy. As these AI strategies evolve, so will the transactions for investing in AI.

2 In some cases, the SPAC transaction will be accomplished by issuing restricted securities to the AI fintech company's stockholders in exchange for the AI fintech company's securities. The SPAC will only file a proxy statement to solicit SPAC stockholder approval of the transaction and will not file a registration statement. The AI fintech company's stockholders will either need registration rights or will rely on the exemption from registration contained in Rule 144 to sell the shares into the market.

The American Data Privacy and Protection Act: Is Federal Regulation of AI Finally on the Horizon?

Niketa K. Patel, Tori K. Shinohara, Jennifer M. Rosa, Arsen Kourinian, Howard Waltzma and Brendan J. Harrington

Summary of Key Points

An omnibus federal privacy bill with significant bipartisan support is currently under congressional review and, if enacted, could dramatically increase oversight of how companies use artificial intelligence ("AI") in their businesses.

This article discusses the bill, which, even if not enacted, provides valuable insights as to potential future regulation of AI.

On July 20, 2022, the House Energy and Commerce Committee approved the proposed American Data Privacy and Protection Act (ADPPA) by a 53-2 margin.[1] The bill would create national standards and safeguards for personal information collected by companies, including protections intended to address potentially discriminatory impacts of algorithms.

Although Congress is unlikely to enact the bill between now and the end of the year, the ADPPA represents progress toward a comprehensive data privacy law in the United States and is part of a growing trend calling for federal regulation of AI.[2] Although several other federal bills addressing algorithmic decision-making have been introduced in recent years, the ADPPA is the first with significant bipartisan support and momentum, and the first to bundle provisions targeting algorithmic accountability and bias with provisions addressing data privacy and security issues.

Scope and Applicability

If enacted, the ADPPA would apply broadly to organizations and businesses operating in the United States. Key definitions in the proposed legislation include those noted below.

Covered entity is defined as an entity that "collects, processes, or transfers **covered data** and is subject to the Federal Trade Commission Act," in addition to nonprofit organizations and common carriers. Though the definition is undeniably broad, the ADPPA identifies several different types of entities with additional obligations or exemptions. For certain obligations, covered entities are divided by "impact" (i.e., annual global revenue and number of data subjects affected by the entity's operations) and "relationship with the data subject" (e.g., direct, third-party, or service provider relationships). By way of example, a "large" entity is defined as one with annual gross

[1] *See* American Data Privacy and Protection Act, H.R. 8152, 117th Cong., https://www.congress.gov/bill/117th-congress/house-bill/8152/text#toc-H4B489C75371741CBAA5F38622BF082DE; American Data Privacy and Protection Act Draft Legislation Section by Section Summary (2022), S. Comm. on Commerce, Science, and Transportation, https://www.commerce.senate.gov/services/files/9BA7EF5C-7554-4DF2-AD05-AD940E2B3E50.

[2] *See* Blueprint For An AI Bill Of Rights: Making Automated Systems Work For The American People, White House Office of Science & Technology Policy, https://www.whitehouse.gov/ostp/.

revenues of at least $250 million and that has collected **covered data** on more than 5 million individuals or devices or has collected **sensitive covered data** of more than 100,000 individuals or devices.

Covered data is defined as "information that identifies or is linked or reasonably linkable to one or more individuals, including derived data and unique identifiers." Importantly, both employee data and publicly available data are excluded from this definition. Certain types of covered data are defined as **sensitive covered data**, which would include government identifiers (such as driver's license or Social Security numbers) as well as "traditionally" sensitive information related to health, geolocation, financials, log-in credentials, race, and sexual history or identity. Sensitive data may also include other categories, such as television viewing data, intimate images, and "information identifying an individual's online activities over time or across third-party websites or online services."

A **service provider** is defined as "a person or entity that collects, processes, or transfers covered data on behalf of, and at the direction of, a covered entity for the purpose of allowing the service provider to perform a service or function on behalf of, and at the direction of, such covered entity." Notably, the ADPPA would place direct obligations on service providers, including obligations not found in state privacy laws such as the prohibition of transferring data, except to another service provider, without affirmative express consent.

A **third-party collecting entity** is defined as "a covered entity whose principal source of revenue is derived from processing or transferring the covered data that the covered entity did not collect directly from the individuals linked or linkable to the covered data." Third-party collecting entities would be required to provide consumers with notice of their activity and register with the Federal Trade Commission ("FTC") if they process data pertaining to more than 5,000 individuals or devices that are reasonably linkable to an individual, as well as provide consumers the opportunity to require such entity delete a consumer's covered data.

Oversight of AI and Algorithmic Decision-Making

With respect to AI, the ADPPA includes a provision—Section 207: Civil Rights and Algorithms—under which covered entities or service providers "may not collect, process, or transfer covered data in a manner that discriminates in or otherwise makes unavailable the equal enjoyment of goods or services on the basis of race, color, national origin, sex, or disability." The two limited exceptions are a covered entity's self-testing to prevent or mitigate unlawful discrimination and a covered entity's efforts to diversify an applicant, participant, or customer pool.

Unlike most existing state privacy laws, Section 207 of the ADPPA would go a step further by requiring companies to evaluate certain artificial intelligence tools and submit those evaluations to the FTC.

Which entities are subject to Section 207?

Covered entities and service providers that develop algorithms to collect, process, or transfer covered data or publicly available information would be required to conduct **algorithm design evaluations** prior to deploying the algorithms in interstate commerce. In addition, any large data holder that uses an algorithm "that may cause potential

harm to an individual," and uses such algorithm to collect, process, or transfer covered data, would also be required to conduct an **algorithm impact assessment** on an annual basis.

What is an "algorithm"? The bill defines a "covered algorithm" as "a computational process that uses machine learning, natural language processing, artificial intelligence techniques, or other computational processing techniques of similar or greater complexity that makes a decision or facilitate human decision-making with respect to covered data, including to determine the provision of products or services or to rank, order, promote, recommend, amplify, or similarly determine the delivery or display of information to an individual." This definition is extremely broad and would cover almost any decision that utilizes automation as part of the decision-making process, even if the ultimate decision is made by a person.

What is an "algorithm design evaluation"? According to the proposed bill, covered entities and service providers must evaluate the design, structure, and data inputs of the algorithm to reduce the risk of potential discriminatory impacts. The draft legislation emphasizes that algorithm design evaluations must occur at the design phase, including any training data used to develop the algorithm. The ADPPA would also require the use of an external, independent researcher or auditor to conduct the evaluation to the extent possible. The covered entity or service provider would be required to submit the evaluation to the FTC no later than 30 days after completion of the evaluation and to make it available to Congress upon request.

What is an "algorithm impact assessment"? For large data holders who use algorithms that may cause potential harm to an individual, and that use such algorithms to collect, process, or transfer covered data, an algorithm impact assessment is also required. The draft bill provides a detailed description of these assessments and requires that they include:

- A detailed description of the design process and methodologies of the algorithm;
- A statement of the algorithm's purpose, its proposed uses, and its foreseeable capabilities outside of the articulated proposed use;
- A detailed description of the data inputs used by the algorithm, including the specific categories of data that will be processed and any data used to train the underlying model;
- A description of the outputs produced by the algorithm;
- An assessment of the necessity and proportionality of the algorithm in relation to its purpose, including the reasons an algorithm is superior to a non-automated decision making process; and
- A detailed description of steps to mitigate potential harms.

Large data holders would be required to submit the impact assessment to the FTC no later than 30 days after completion of the assessment and continue to produce assessments on an annual basis. As with algorithm design evaluations, the proposed legislation would require the use of an external, independent researcher or auditor to conduct the algorithm impact assessment, to the extent possible.

The level of prescriptive detail may require many companies, and especially large data holders, to dedicate significant resources to assessing their

algorithmic tools during the development phase and additional resources to monitoring those same tools during and after development.

Which "potential harms" require an algorithm impact assessment? The following potential harms are expressly highlighted in the text of the bill, suggesting that these are areas of focus for lawmakers:

(i) Potential harms related to individuals under the age of 17;

(ii) Potential harms related to advertising for, access to, or restrictions on the use of housing, education, employment, healthcare, insurance, or credit opportunities;

(iii) Potential harms related to determining access to, or restrictions on the use of, any place of public accommodation, particularly as such harms relate to protected characteristics, including race, color, religion, national origin, sex, or disability; and

(iv) Potential harms related to disparate impact on the basis of individuals' race, color, religion, national origin, sex, or disability status.

The language of the proposed bill suggests that this list of potential harms is not exhaustive. It is also worth noting that the bill is under consideration at a time when there is significant regulatory attention on ad targeting and digital marketing, including by the Consumer Financial Protection Bureau, which recently issued an interpretive rule on digital marketing and expressed concern over discriminatory conduct online and "digital redlining."[3]

What does it mean to "discriminate" under Section 207? One of the key questions raised by the proposed legislation, and one that would be critical to assessing compliance, is what exactly does it mean to "discriminate" under Section 207 of the ADPPA? While Section 207's reporting requirements involve descriptions of any "disparate impact" resulting from the deployment of an algorithm in a covered entity's business practices, it is unclear what legal standards would be used in assessing discrimination or disparate impact under the proposed legislation and what type of business justification might suffice to satisfy the proposed bill's requirements. Depending on the algorithm, it may be very difficult—if not impossible—to completely eliminate all disparate impact against any protected classes, even when using objective and facially non-discriminatory criteria. In addition, the proposed legislation refers to "protected characteristics," but this term is not defined, nor does the proposed legislation reference any federal or state anti-discrimination laws that explicitly enunciate the so-called "prohibited bases" that such laws are designed to protect. Moreover, the proposed bill does not address how companies are expected to perform testing in the absence of demographic data such as race or national origin and whether proxying methodologies (such as the Bayesian Improved Surname Geocoding—or "BISG") would be required.

Enforcement

The ADPPA would create a Bureau of Privacy at the FTC to enforce its provisions, and any ADPPA violation would be treated as a violation of a rule defining an unfair or deceptive act or practice

[3] *See* US CFPB Takes Aim at Digital Marketing Providers with New Interpretative Rule (https://www.mayerbrown.com/en/perspectives-events/publications/2022/08/us-cfpb-takes-aim-at-digital-marketing-providers-with-new-interpretative-rule).

("UDAP") under section 18(a)(1)(B) of the Federal Trade Commission Act (15 U.S.C. 57a(a)(1)(B)).

With respect to Section 207, the ADPPA would authorize the FTC to promulgate regulations to establish processes by which large data holders can submit impact assessments and exclude from assessment "any algorithm that presents low or minimal risk for potential for harms to individuals." The ADPPA would also require the FTC to publish guidance within two years of the bill's enactment regarding compliance with Section 207 and a study within three years of the best practices for assessment and evaluation of algorithms and methods to reduce the risk of harm. These publications may help provide guidance to companies as they navigate compliance and dedicate resources to the evaluation of algorithmic tools.

Although the ADPPA as drafted includes a private right of action about which a number of business groups have raised concerns, it, importantly, would *not* apply to Section 207's provisions related to potential discrimination. Instead, the FTC and state attorneys general would be empowered with enforcement authority with respect to Section 207.

What's Next?

Despite the bipartisan support, the bill has faced significant resistance from California lawmakers who argue that the bill would preempt the California Privacy Rights Act ("CPRA"), which they argue offers stronger protections to California residents (though a number of experts, including a former Chairman of the FTC, have questioned whether the CPRA actually provides stronger protections). Several state attorneys general have also sent a joint letter to Congress expressing the urgent need to amend the bill to explicitly allow states to pass potentially more expansive privacy, data, and artificial intelligence-related requirements in the future as technology and online practices evolve. Conversely, a number of business groups have expressed concerns that the bill does not effectively preempt state laws, leaving in place, at least to a certain extent, a patchwork of privacy laws across the United States.

Even if its enactment is unclear, the ADPPA provides significant insights as to the type of oversight of AI tools that lawmakers and regulators may seek to exercise in the near future. It is an issue that is likely to receive continued focus by the federal government, as demonstrated by the White House Office of Science & Technology Policy's recent unveiling of a Blueprint for an AI Bill of Rights.

Companies may wish to consider developing internal impact assessment forms for design teams to fill out during the development phase of algorithmic products, paying particular attention to data integrity and data inputs; human oversight, monitoring, and control; and, potentially, disparate impact analyses. These impact assessment forms and related processes could be embedded into existing governance protocols, and training could be arranged for relevant stakeholders. Companies may also consider whether their organizations would benefit from the addition of an AI committee or whether existing risk committees or other bodies can expand their remit to assess impacts of algorithmic applications. The teams conducting the impact assessment would benefit from being cross-functional and diverse—design and technology experts, risk and/or compliance strategists, marketing professionals, ethicists, and lawyers can all be important advisors during this process.

Explainability Is an AI Tool's Best Defense

Christopher Leach, Alex Lakatos, Reginald Goeke

Summary of Key Points

- Explainability—the ability of human beings within your organization to explain an artificial intelligence and machine learning ("AI/ML") tool—is essential for your company to successfully navigate likely litigation related to that tool.
- Explaining AI/ML tools is difficult because these tools handle copious amounts of data, which could include sensitive information; change their functionality in response to what they "learn" over time; are constantly updated by many humans who may not document their code for laypeople (if they document at all); and produce output that requires special expertise to interpret.
- Plaintiffs' lawyers may weaponize the inherent difficulties of AI/ML tools by, for example, (a) making the preservation of your AI/ML tool's code documentation look as if it were simple and then asking for sanctions based on an inference that the "missing" data was not preserved in bad faith; (b) asking for sensitive information and then making your productions seem inadequate and requesting their experts get direct access to the tool; and (c) having laid the groundwork through those tactics, retaining experts to challenge your AI/ML evidence's authenticity and filing motions *in limine* to exclude your company's use of this evidence.
- Improving explainability will better prepare your company to successfully navigate litigation related to the AI/ML tool.

Financial services firms are increasingly employing AI/ML for an expanding number of uses. While these tools can be accretive to business, the decisions they assist with likely will lead to litigation in some way—a fair lending suit, employment issues associated with pay, disputes with customers or regulators regarding the handling of customer accounts, or, for companies that design these tools, disputes with customers in the event that something goes wrong.

All this can be a huge headache in litigation, leading to discovery expenses, business leaders occupied with litigation, and (heaven forbid) adverse judgments—if companies don't design their programs right.

But with some foresight, companies can set themselves up not only to benefit from the AI/ML tools but also to avoid costly and annoying pitfalls if things end up in court. Of course, nothing is foolproof, so companies also should be aware of the potential tactics that opposing counsel might employ to gum up your case.

In this article, we make the case that the ability to explain an AI/ML tool is essential for a company to navigate litigation related to that tool successfully. After laying out what we mean by "explainability," we then set out the various aspects of AI/ML tools that make explainability so difficult, followed by some specific litigation-based examples of how these difficulties arise. We conclude with some tips for how to structure your business and your product to ensure that

when a dispute arises, your business is ready and not left scrambling.

What is "explainability"?

One of the core concepts in creating a litigation-ready AI/ML tool is explainability. The reason for this is obvious: the AI/ML tool is incredibly complicated and cannot testify, sign an affidavit, walk a regulator through a particular practice, or explain an action to a customer (to avoid disputes in the first place). Your company's ability to justify its actions to courts, adversaries, and customers depends in large part on whether human beings within your organization can articulate what happened and why.

Explainability, in general terms, has three components:

- Transparency: easy identification of the important factors in the tool's operation;
- Interpretability: easy identification and explanation of how the tool weighs those factors and derives them from its input data; and
- Provenance: easy identification of where input data originates and what the data contains.

These principles arm companies with the ability to frame the discussion in terms of decisions people made, not results that the AI/ML tool shot out. In other words, by focusing on explainability, companies can ensure that their AI/ML tool *is a tool* and not a decisionmaker.

WHAT MAKES EXPLAINING AI/ML SO DIFFICULT?

What makes AI/ML explainability so difficult is that, frankly, there is a lot going on. What goes in, what comes out, and what happens in between involves massive, changing data engaging with computer programming that evolves its processes as it consumes additional data. These features—and others—not only complicate explainability but also are among the key reasons why litigating cases involving AI/ML tools can be tricky, especially in discovery.

Voluminous Data. A feature of AI/ML tools is that they consume copious amounts of data. This is, in part, because these tools improve as operators "train" the AI/ML tool with more data. Indeed, AI often functions by analyzing all the data that is available, e.g., reviewing all transactions, customer data, behavioral data, and the like to spot money laundering risks or to assess creditworthiness. Producing and reviewing this data, as litigation often requires, poses significant challenges.[1] Moreover, as algorithms become more sophisticated, they require even greater amounts of data.

Sensitive Data and Trade Secrets. Depending on the tool, data being fed into an AI/ML tool could include sensitive personal information, financial information, or even health and spending records. Keeping this information secure obviously induces some tension with respect to obtaining maximal understanding of the tool's inner workings. That goes triple for the underlining intellectual property of the tool itself, which often reflects valuable trade secrets. And this sensitivity also presents tricky issues in litigation when plaintiffs ask for the underlying data.

A Moving Target. Good AI/ML tools are not static. As the name suggests, AI/ML tools "learn," change, and improve functionality over time as

[1] Artificial Intelligence and Privacy, Datatilsynet (Norwegian Data Protection Authority) at page 4 (January 2018), available at https://www.datatilsynet.no/globalassets/global/english/ai-and-privacy.pdf.

they ingest data from newer and more varied sources. A machine learning system that recommended extending credit on one day might make a different recommendation on a later day based on the system having seen more data, having learned more, and having refined its internal model in the interim. This presents discovery challenges for data. For example, is it even possible to go back and identify the data that the machine learning system trained on at a particular moment? Indeed, because many programs overwrite their code as they evolve, obtaining a historical understanding of your data—either to explain or to produce because of retention obligations—is difficult to do after the fact. The AI system, however, may not be configured to retain values that change or are overwritten as the AI learns.

The "Black Box." In the days of yore, algorithms were rules-based decisionmakers and could be understood by studying those rules and the underlying computer code. For example, a rules-based program might have a rule that provides that if a borrower has a certain debt-to-income ratio above a certain amount, then the lender will not extend any additional credit to that particular borrower. AI/ML tools are not that simple and often lack deterministic rules. Rather, these tools use statistical models and probabilistic rules. Machine learning might approach the problem above by building a model to answer the question: how much does this potential borrower resemble those who have paid as agreed versus those who have defaulted?

But how did the AI/ML tool actually reach that result? Although you may know the data inputs, it is harder to know what it considered, how much weight it gave to any particular factor, and how the factors might be interrelated. The answers may be unintuitive, which, indeed, is the point of using an AI/ML tool rather than humans. As a result, knowing the output may not provide any insight into how the AI/ML tool arrived at its decision. Indeed, because AI/ML tools are only concerned with the specific outcomes that their engineers instruct them to care about, the tools may take a path to get to an end point that humans would consider to be cheating, undesirable, or otherwise inconsistent with their intentions.

Development (By Humans). Good AI/ML tools not only update themselves but also are often in a process of constant updating and revisions by software engineers and data scientists, who may be doing anything from experimenting with new techniques for analysis to tweaking the inputs or outputs. There may not be one static set of code to produce but millions of lines, with hundreds or thousands of owners, in a constant state of flux. Documenting these changes in real time obviously is more helpful when it comes to explainability. But computer scientists and software engineers, particularly those in nimble fintech startups, may not have a strong culture of documentation, and, in any event, that documentation may be more geared toward the needs of other computer scientists, not of litigants. And even then, programmer shorthand can easily be taken out of context by litigation adversaries.

That's just for existing systems. What happens when newer AI/ML tools replace older ones? The older system often is not maintained in a useable form. Indeed, when an AI/ML tool is shelved, incentives to maintain that obsolete system fade.

The Output. Beyond trying to understand the inputs, the black box, and changes over time with both, there is an issue with the output from the model. Like the incoming data, the outputs can be huge, complex, and evolving themselves,

requiring special expertise to read and interpret. And even accessing the output from a technical level can be challenging. For example, the information that AI system outputs may be (1) stored in deep storage, so that information first must be moved to fast storage before it can be searched; (2) stored in a proprietary format as opposed to commonly known formats such as .xls or .csv; or (3) subject to search and review only using specialized tools that may exist only in-house and that may be understood only by in-house engineers.

HOW COULD THESE CHALLENGES EMERGE IN LITIGATION?

These challenges associated with an AI/ML tool can emerge in various forms during the course of litigation. Below, we discuss several areas where companies can run into real trouble when the complexities of modern AI/ML tools interact with the rules of civil procedure and plaintiffs' lawyers who may try to use the inherent difficulties of AI/ML tools to their advantage. Although these issues could come up in countless iterations, we will discuss here issues related to preservation of documents, production of documents, and authentication at trial.

Preservation. Parties anticipating or conducting litigation have obligations under the Federal Rules of Civil Procedure to "take reasonable steps to preserve" electronically stored information ("ESI").[2] This likely would include ESI related to a company's AI/ML tools, in the appropriate case.

Here, the key risk of failing to preserve documents would be sanctions for "spoliation." Spoliation involves the destruction of, significant alteration of, or failure to properly preserve evidence in pending or reasonably foreseeable litigation. When courts determine that a party has wrongfully failed to preserve evidence, courts have broad discretion to impose sanctions, which could include, among other things, precluding a party from presenting evidence on topics addressed in evidence that was subject to spoliation, allowing evidentiary inferences that the missing evidence would have been adverse to the party that failed to preserve it, finding certain issues conclusively established against the party that failed to preserve evidence, and entering a default judgment against the party responsible for the spoliation.[3] These are serious sanctions, and courts generally consider whether the party's preservation efforts were reasonable and undertaken in good faith when deciding whether/which sanctions to impose.

Preserving information related to AI/ML tools may seem like a daunting exercise given the difficulties outlined above. But the drafters of the Federal Rules of Civil Procedure imposed some important guardrails in 2015 to protect companies from crushing sanctions for non-culpable spoliation. Specifically, under Rule 37(e), sanctions are only appropriate if the lost information cannot be restored or replaced through additional discovery and may be no greater than necessary to cure the prejudice caused by the loss of information. Further, under Rule 37(e), a court must conclude that a party's spoliation of ESI was intentional (and not merely negligent or grossly negligent) before imposing more serious discovery sanctions such as an adverse inference or default judgment. This rule helps protect parties using AI/ML tools, although some courts have been willing to infer intent

[2] Fed. R. Civ. P. 37(e).

[3] See, e.g., FRCP 37; NY CPLR § 3126.

from the circumstances of the preservation failure itself,[4] and some courts have allowed the jury to decide in the first instance whether the failure to preserve ESI was intentional.[5]

With this background law, it is important to contemplate how creative plaintiffs' lawyers would seek to exploit the situation. Expect plaintiffs first to create a record regarding preservation demands through a series of letters and meet-and-confers making unreasonable demands, detailing the ways in which the company has "fallen short." Then they might seek discovery of documents (such as manuals, code and code documentation) and take depositions (e.g., of corporate representatives, of engineers) to test what preservation might have been feasible. They also might hire "experts" in connection with a court challenge to take unrealistic positions on what preservation would have been possible. Then, before the court, they could make the dispute appear as simple as possible, asserting that preservation is simple and inexpensive to achieve. If successful, the plaintiffs will ask for sanctions, likely based on an inference that the "missing" data was not preserved in bad faith.

Production. Generally, parties must produce relevant information as part of the discovery process, typically in response to requests for production from the other side. These discovery requests can include all aspects of the AI/ML tool—the data inputs, the outputs, and the source code itself. Of course, just because plaintiffs ask does not mean you need to produce. The production obligation is not unlimited, and (at least in federal court) document requests must be proportional to the needs of the case.[6] Similarly, state courts generally permit objections to discovery requests that are unduly burdensome. In disputes over production of AI, however, there is little guidance over where the proportionality/unduly burdensome line should be drawn.

If discussions are not fruitful, then the parties can seek guidance from the court—either as a motion to compel production or a motion for a protective order. In that proceeding, the court would consider the relevance of the information and the burdens of production, with the (somewhat obvious) rule of thumb that there is a greater chance of compelled production where the data is central to the case. So, for example, courts have required production of source code in patent or copyright disputes where the that information was essential to determine the scope of the product at issue but have been less inclined to order production in cases where the merits are more tangential to the AI/ML tool, such as in certain false advertising suits. Note that the mere fact that source code may be (or reflect) sensitive trade secrets often is not sufficient to preclude production, as courts often view the protective order in the case as sufficient to safeguard commercially sensitive material. Once the court has determined and entered an order governing the required scope of discovery, failure to comply may lead to discovery sanctions such as those described above.[7]

---

[4] See, e.g., *O'Berry v. Turner*, 2016 WL 1700403 (M.D. Ga. Apr. 27, 2016) (holding that failure to preserve ESI, reliance on a single hard copy, and loss of that hard copy supported a finding of intent to deprive).

[5] See, e.g., *Cahill v. Dart*, 2016 WL 7034139 (N.D. Ill. Dec. 2, 2016) (holding that the jury should make the decision whether prison officials had intentionally allowed a crucial part of a videotape segment to be overwritten).

[6] See FRCP 26(b)(1).

[7] See, e.g., FRCP 37(b)(2).

So, how do plaintiffs' lawyers exploit the difficulties of an AI/ML system? As with preservation issues, this often starts with broad requests that seek sensitive information, including trade secrets and sensitive customer information. These requests are then often followed by letters memorializing plaintiffs' view that any productions were inadequate and requesting that their experts or document specialists obtain direct access to the AI/ML tool. And then, when that does not work, they will seek a motion before the court, backed up by sanctions.

Admissibility at Trial. If the case makes it to trial, companies defending suits involving AI/ML tools will want to present evidence explaining how it works. One possible hurdle here is authentication—the evidentiary requirement that the party presenting the evidence must demonstrate the evidence is what it purports to be. When determining whether a party has laid proper foundation for admission of computer-generated evidence, courts consider, among other things, whether the computer was standard and in good working order, whether the operators of the equipment were qualified, whether proper procedures were followed, whether reliable software was used, whether the program operated properly, and the exhibit derived from the computer. This standard is flexible and often more complicated to show as the complexity increases. Generally, the considerations include (1) the quality of the data input, (2) the complexity of the algorithm, (3) whether the problem is routine or novel, and (4) whether the output can be tested and verified. A 2017 amendment to the federal rules related to "a record generated by an electronic process or system" conceivably could help make things easier, but that rule is intended for routine computer-generated evidence such as an electronic phone log. By contrast, AI models with inputs, weights, and outputs that are in flux, or that are novel and hard to comprehend, may encounter authentication challenges.

Plaintiffs trying to muck up trial usage of your AI/ML model would likely have laid the groundwork through the discovery tactics outlined above, including seeking overreaching discovery, and then retain experts to challenge authenticity and file motions *in limine* to exclude your company's use of AI/ML evidence.

Of course, having your AI/ML tool excluded from evidence at trial could be devastating to your litigation case. For example, if your company relies on immunity provided by Section 230 of the Communications Decency Act,[8] understanding exactly how your AI/ML tool works could be essential to proving that your company was not a content creator. Or if your company uses an AI/ML tool in connection with customer service calls, explaining how the tool works could be essential in defending against claims that, for example, your tool can identify individual voices sufficient to trigger liability under biometric privacy statutes.

HOW CAN YOU MITIGATE LITIGATION RISKS FOR AI/ML TOOLS?

Now that we see all the ways in which plaintiffs can weaponize the complexity of AI/ML tools to derail corporate defense efforts, what can companies do to ready themselves for potential litigation?

[8] 47 U.S.C. § 230 (c)(1) ("No provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider.").

Improving Explainability. Improving explainability is one of the essential pieces. As we previewed above, one of the key goals of explainability is to make clear that business decisions are being made by people and that the AI/ML tool remains just that. Doing so requires thoughtful engagement throughout not only the company but also the lifecycle of the AI/ML tool. We recommend these steps:

- The management team makes clear up front how the company wants the AI/ML tool to work, recognizing that specifications are not simply technocratic decisions left for the engineers and data scientists but are, in fact, business decisions;
- The company's e-discovery/information governance team specifies to data scientists, computer scientists, software engineers, and technicians how the company wants to store and access input, outputs, change logs, models, and the like;
- Legal and compliance are involved in these discussions early and throughout the process;
- The company regularly tests and modifies the AI/ML tool to keep it working as the management team and the e-discovery/information governance team intended. To that end, companies should consider hiring or designating personnel as "AI Sustainers" whose primary responsibility is testing and modification; and
- The company has on-hand individuals who can explain the tool's results.

In addition to those steps—which concern human supervision of the AI/ML tool—companies could further explainability goals by including features to the tool, such as:

- Code that permits auditing and testing;
- Employing "Explainable AI," a service offered by some AI/ML companies that provides a better window into "black box" decision-making;
- Extra documentation that explains how the AI/ML tool works and what choices were made about its features and functionality, for the benefit of current in-house employees, later in-house employees, and later retained experts; and
- Thoughtful decisions about which facts and data to preserve and which to overwrite.

Drawing on Explainability in Litigation. Explainability will be invaluable when confronted with the problems of production, preservation and proof described above.

Preservation. Explainability affords the company several advantages in connection with its obligation to preserve documents. The documentation described above creates records in advance that the company knows are important and hopefully allows more targeted retention and preservation processes. And having gone through that process provides the company with a rationale to defend its preservation choices.

Production. These steps obviously make production easier because the company has pre-identified the documents most likely to be relevant, with ready-made explanations against expansive and burdensome requests for additional data of marginal relevance. But going further, working on explainable AI/ML outputs will ensure that exported data is both comprehensible and portable for production to plaintiffs and the court.

Proof and Authentication. Finally, if your company can explain to itself the inputs, outputs, and processes of the AI/ML tool, your company

will be in a good position to explain to a court how your AI/ML tool works convince a court of these issues.

Indeed, the ultimate goal of implementing an explainable AI/ML tool is to short-circuit the inevitable "battle of the experts" throughout the process, including at trial. Without an explainable tool, plaintiffs will provide expert testimony stating that the AI should have resulted in a set of decisions that, not surprisingly, establish that they were harmed, with the defendants relying on their own expert for the opposite. This scenario allows plaintiffs' experts to provide the court with simple assertions rebutted only by the competing expert rather than the "fact" of the AI/ML tool itself. This all inures to the plaintiffs' benefit because fact finders often default to simpler explanations.

WHAT DOES THIS ALL MEAN?

With AI tools ubiquitous in financial services, they too will be ubiquitous in financial services litigation. This is unfortunate for the defendants of those suits, who obviously would rather not spend the time and money needed to litigate. But thinking critically about how to get your AI/ML tool ready for litigation should dramatically reduce the expense incurred when companies are forced to do so on more constrained litigation timelines.

US Insurance Regulators' Evolving Perspectives on Artificial Intelligence

Paul Chen, Vikram Sidhu and Yuliya Feldman

Summary of Key Points

- The National Association of Insurance Commissioners ("NAIC") has a committee and working groups considering the use of big data and artificial intelligence ("AI") in the industry and evaluating existing regulatory frameworks for their use. In addition, the NAIC has a forum for ongoing discussion among insurance industry stakeholders around these issues.
- These NAIC initiatives could lead to the development of or modifications to model laws, regulations, handbooks, and regulatory guidance.
- Some state insurance regulators, such as the New York Department of Financial Services, the California Department of Insurance, and the Connecticut Insurance Department, have issued circular letters and bulletins highlighting their concerns about bias and discrimination resulting from the use of AI and machine learning ("ML") in insurance.
- Colorado has enacted a statute that requires its insurance commissioner to adopt rules prohibiting insurers from using algorithms or predictive models that use external consumer data and information sources in a way that unfairly discriminates. Other states have, or have had, similar legislation pending.

As technological innovation has gathered speed in the insurance industry over the past decade, state insurance regulators have tried to enable the implementation of insurance technology while balancing consumer protection concerns. While recognizing that advancements in insurtech certainly enable the delivery of a broader range of insurance products through streamlined underwriting processes and the payment of claims more efficiently through more effective data analytics, state insurance regulators continue to be concerned about ensuring that consumers understand the insurance products that they are buying, that insurance products are accessible and fairly priced without reference to criteria that could be regarded as discriminatory, and that individual consumer data is adequately protected and kept private. Among the key areas on which state insurance regulators have been focusing their attention with respect to innovation and technology in insurance is the use of artificial intelligence ("AI"), including machine learning ("ML"), in the insurance industry.

In the United States, state insurance regulators' efforts with respect to studying, assessing and potentially regulating the use of AI has been led principally by the National Association of Insurance Commissioners ("NAIC"), which is the association of the US insurance regulators from all 50 states, DC and the territories. In addition, certain US states have taken the lead individually in assessing the potential regulatory considerations with respect to the use of AI in insurance.

NAIC

Among the 2022 priorities for the NAIC is to analyze AI advancements to assess if current state laws and regulatory tools are sufficiently protecting consumers. This work is centralized within the NAIC's Innovation, Cybersecurity, and Technology (H) Committee (the "ICT Committee"). Although this committee has a broad mandate with respect to innovation, cybersecurity, privacy, e-commerce and technology in insurance, one of its key working groups is the Big Data and Artificial Intelligence (H) Working Group (the "BD/AI Working Group").

The BD/AI Working Group is tasked with, among other things, researching the "use of big data and [AI] including [ML] in the business of insurance and evaluat[ing] existing regulatory frameworks for overseeing and monitoring their use"; "[reviewing] current audit and certification programs and/or frameworks that could be used to oversee insurers' use of consumer and non-insurance data, and models using intelligent algorithms, including AI"; and "[assessing] data and regulatory tools needed for state insurance regulators to appropriately monitor the marketplace, and evaluate the use of big data, algorithms, and machine learning, including AI/ML in underwriting, rating, claims and marketing practices."

The BD/AI Working Group met on August 10, 2022 at the NAIC Summer 2022 National Meeting. At the meeting, the working group received an analysis of the results of an AI/ML survey for the private passenger auto line of business, which was done in [2021]. There is an AI/ML survey being developed for the home line of business, which is in the final stages of development; once the NAIC programs the survey into its systems, 10 states will formally issue the market conduct data call to insurers. Finally, an AI/ML survey for the life line of business is in the development phase.

In addition, the BD/AI Working Group has a "Third-Party Data and Model Vendors workstream." The workstream is considering several potential initial steps for enhanced regulatory oversight of third-party data and model vendors, including requiring contracting insurers to certify that the models that are being used comply with certain standards and developing a library of third-party vendors.

At the NAIC Summer 2022 National Meeting, the ICT Committee held a meeting of the Collaboration Forum on Algorithmic Bias, which was established by the NAIC earlier in 2022 as a platform for multiple NAIC committees to work together to identify and address foundational issues and develop a common framework that can inform the specific workstreams in each group. Rather than being a single event, the Collaboration Forum is intended to promote ongoing discussion among insurance industry stakeholders during regularly hosted events and presentations. The Collaboration Forum on Algorithmic Bias was designed to cover issues such as what kinds of algorithms raise concerns for insurance regulators, how might bias arise in algorithms, which tools might be effective in minimizing bias and detecting bias, and what are potential regulatory frameworks for addressing algorithmic bias.

The presentations made during the Collaboration Forum at the Summer 2022 National Meeting covered the following topics: Perspectives on AI Risk Management and Governance, Bias Detection Methods and Tools, Ethical and Responsible Use of Data and Predictive Models, Today's Approaches to Algorithmic Bias, and Risk of Biased AI. Some of

the key themes explored during these presentations were the following:

- *Risk Management Approach to AI:* Several presenters discussed that, in the absence of more specific guidance from insurance regulators on the use of AI/ML, the industry should treat its use of AI/ML as part of regular risk management. That is, a comprehensive AI/ML risk management and governance framework should include the following components: development and communication of written policies and procedures (including assignment of responsibility and accountability with respect to such policies and procedures), training and monitoring with regard to the policies and procedures, and taking corrective action (and documenting that action) when the policies and procedures are not followed.
- *Ethical Use of Data and Predictive Models:* Several presenters discussed the principles that they believe should guide the industry's use of AI/ML, including fairness, safety, transparency and accountability. There was significant discussion of how the industry, guided by these principles, could avoid bias in all stages of AI/ML model development, including during the pre-design, design and development, testing and evaluation, and deployment stages.
- *The Need for Testing:* Several presenters emphasized the need for testing as a critical tool for identifying unintended discrimination. There are several forms of testing available that could be used to identify bias, including the Control Variable Test, the Interaction Test, the Nonparametric Matching (Matched Pairs) Test, and the Double Lift Chart. According to the presenters, the appropriate test for any particular model will vary based on the model type, the intended use, the output, the volume of data available, and the granularity of protected class data available.
- *Access to Protected Class Data:* The issue that insurers currently do not have systematic data about policyholders' membership in protected classes was raised several times during the discussion. The lack of this data could make testing for bias more difficult.
- *The Need for Diversity:* Several presenters highlighted the importance of diversity in combating algorithmic bias. They explained that to prevent bias in the development stage, models should be established with diverse users in mind, and a diverse and inclusive workforce is critical for the oversight or monitoring of AI/ML use because diverse perspectives can help identify bias.
- *Model Explainability:* Several presenters emphasized the importance of transparency and model explainability. In furtherance of this guiding principle, a proposal was made to develop model cards, which would present certain basic information about an AI/ML model (e.g., a description of the model goals, limitations of the model, trade-offs with respect to the use of the model and performance of the model). This proposal was described as being the equivalent of nutrition labels for AI/ML models.

The insights shared at the Collaboration Forum will be used by the ICT Committee and its BD/AI Working Group to evaluate existing regulatory frameworks for overseeing and monitoring the use of big data, algorithms, and machine learning—including AI/ML in underwriting, rating, claims, and marketing practices of insurers—potentially leading to the development of or modifications to model laws, regulations, handbooks and regulatory guidance.

State Developments

In addition to the work being done on the use of AI in insurance at the NAIC, several states have also issued guidance to the insurance industry with respect to the use of AI, including big data and ML. For example, the New York Department of Financial Services ("NY DFS") issued its Circular Letter No. 1 (January 18, 2019), which resulted from an investigation into New York life insurers' underwriting guidelines and practices. To address concerns about potential unlawful discrimination, the circular letter set forth two guiding principles for New York insurers that use external data in underwriting: (i) that insurers using external data sources must independently confirm that the data sources do not collect or use prohibited criteria; and (ii) that insurers should not use external data unless they can establish that it is not "unfairly discriminatory" in violation of applicable law—i.e., using external data only if the insurers are confident that the use of the data is demonstrably predictive of mortality risk and that they can explain how and why this is the case. The circular letter highlighted that NY DFS, like other regulators, continues to be concerned about unlawful discrimination and transparency in the use of data as well as AI and ML in insurance.

Based on similar concerns, the California Department of Insurance ("CDI") recently issued its Bulletin 2022-5 on June 30, 2022. The focus of the bulletin was to address allegations of racial bias and discrimination in marketing, rating, underwriting, and claims practices by insurance companies and other licensees. CDI, like NY DFS, also highlighted concerns about transparency, and noted that the "greater use by the insurance industry of artificial intelligence, algorithms, and other data collection models have resulted in an increase in consumer complaints relating to unfair discrimination in California and elsewhere" and that the "use of these models and data often lack a sufficient actuarial nexus to the risk of loss and have the potential to have an unfairly discriminatory impact on consumers." CDI emphasized in the bulletin that insurers and other licensees must "avoid both conscious and unconscious bias or discrimination that can and often does result from the use of artificial intelligence, as well as other forms of 'Big Data' (i.e., extremely large data sets analyzed to reveal patterns and trends) when marketing, rating, underwriting, processing claims, or investigating suspected fraud relating to any insurance transaction that impacts California residents, businesses, and policyholders." Further, the bulletin provided that "before utilizing any data collection method, fraud algorithm, rating/underwriting or marketing tool, insurers and licensees must conduct their own due diligence to ensure full compliance with all applicable laws."

Similarly, the Connecticut Insurance Department ("CID") issued a bulletin on April 20, 2022 regarding The Usage of Big Data and Avoidance of Discriminatory Practices (which updated and amended a bulletin issued on April 8, 2021). CID highlighted similar themes as its counterparts in New York and California—that insurance companies and other licensees must use technology and data in full compliance with anti-discrimination laws. CID also began requiring a "data certification" that insurance licensees' use of data complies with CID's bulletin and applicable laws; the first certification was due on September 1, 2022.

Some states are taking more robust action and introducing legislation to specifically prohibit discrimination in the insurance industry's use of AI. In July 2021, Colorado enacted a new statute

that requires the Colorado Insurance Commissioner to adopt rules prohibiting insurers from using any external consumer data, information sources, algorithms or predictive models that use external consumer data and information sources in a way that unfairly discriminates based on race, color, national or ethnic origin, religion, sex, sexual orientation, disability, gender identity or gender expression. The Colorado Division of Insurance has conducted several stakeholder meetings to discuss related issues before the Division proceeds with adopting rules on how insurers should test and demonstrate to the Division that their use of big data is not unfairly discriminating against consumers. Other states have, or have had, similar legislation pending.

* * * *

As the use of AI by the insurance industry, including data that feeds into AI and ML, continues to grow, the developments at both the NAIC and at the state level are expected to continue to evolve as well.

UK Government proposes a new approach to regulating artificial intelligence (AI)

Authors: Oliver Yaros, Ondrej Hajda, Mark Prinsley, Valerie Vanryckeghem, Reece Randall and Ellen Hepworth

Other Contacts: David Simon, Ana Bruder and Ulrich Worm

Tags: Technology & IP Transactions, Cybersecurity and Data Privacy, FSRE, Technology, Brexit

The UK Government published the AI Regulation Policy Paper on 18 July 2022. The Policy Paper sets out the Government's vision for the future "pro-innovation" and "context-specific" AI regulatory regime in the UK.

The Policy Paper outlines six cross-sectoral AI governance principles and confirms that the UK Government is not currently planning to introduce new legislation in the UK to regulate AI. However, the UK Government plans to ask existing regulators to interpret and implement the cross-sectoral principles that will be at the heart of UK's new AI regulatory regime. The Policy Paper forms part of the UK Government's National AI Strategy[1] and its AI Action Plan[2].

Organisations that use or sell AI in the UK should monitor the upcoming AI White Paper (expected in late 2022) and announcements from the relevant regulators. Businesses should also consider who is responsible for AI governance and risk management strategy within their organisation, and prepare the align of their internal AI strategy with the proposed principles.

Principles

The Policy Paper presents an early proposal for six cross-sectoral principles that the UK Government is planning to ask regulators to apply in their sector or domain:

1. **Ensure that AI is used safely**: Safety is likely to be a core consideration in certain sectors (such as healthcare or critical infrastructure). However, the Policy Paper suggests that all regulators should take a context-based approach when determining the likelihood of AI posing a risk to safety and take a proportionate approach to managing this risk.
2. **Ensure that AI is technically secure and functions as designed**: AI systems should be technically secure and work as they claim and intend to do. The Policy Paper envisages that functioning, resilience and security of AI systems are tested (subject to context and proportionality considerations) and regulators set out the regulatory expectations in their relevant sector or domain.
3. **Make sure that AI is appropriately transparent and explainable**: The Policy Paper acknowledges that AI systems cannot always be meaningfully explained and in most

[1] https://www.gov.uk/government/publications/national-ai-strategy/national-ai-strategy-html-version

[2] https://www.gov.uk/government/publications/national-ai-strategy-ai-action-plan/national-ai-strategy-ai-action-plan

situations this is unlikely to pose substantial risk. However, the Policy Paper suggests that in certain high-risk settings, decisions that cannot be meaningfully explained might be prohibited by the relevant regulator (for example, a tribunal decision where the lack of explainability would deprive the individual of a right to challenge the tribunal's decision).

4. **Embed considerations of fairness into AI**: The Policy Paper proposes that regulators will define "fairness" in their sector or domain and outline when fairness considerations are relevant (for example, in the case of job applications).

1. **Define legal persons' responsibility for AI governance**: The Policy Paper confirms that accountability for the outcomes produced by AI systems and legal liability must always rest with an identified or identifiable legal person.

2. **Clarify routes to redress or contestability**: According to the Policy Paper, the use of AI should not remove the right to contest a decision where such right is available to individuals and groups outside the AI setting. Therefore, the UK Government will expect regulators to ensure that outcomes of AI systems can be contested in "relevant regulated situations".

The proposed principles build on the five OECD AI Principles[3] and highlight the areas where the UK Government sees the most risk when using AI.

The Policy Paper also confirms that the UK Government will ask the regulators to focus on high-risk concerns (rather than hypothetical or low risks associated with the use of AI) and to consider lighter touch options for regulation (such as issuing guidance or encouraging voluntary measures).

Regulators

The Policy Paper identified the Information Commissioner's Office (ICO), Competition and Markets Authority (CMA), Ofcom, Medicine and Healthcare Regulatory Authority (MHRA), and Equality and Human Rights Commission (EHRC) as the key regulators in its new regime.

While many UK regulators[4] and UK Government departments[5] have already started to take action to support the responsible use of AI, the Policy Paper highlights some of the current challenges faced by businesses, including a lack of clarity, overlaps, and inconsistency between different regulators.

The risk of multiple regulators being asked to interpret and enforce a set of common principles is that businesses will be given inconsistent or contradictory guidance or guidance which leads to duplication of efforts. The Policy Paper acknowledged this risk and stressed that the UK

[3] The OECD Council adopted the OECD AI Principles to promote use of AI that is innovative and trustworthy and that respects human rights and democratic values in May 2019.

[4] For example, the ICO published *Guidance on AI and Data Protection* and *Guidance on Explaining decisions made with AI*, the Bank of England and FCA published the *AI Public-Private Forum Final Report*, the FCA commissioned The Alan Turing Institute to publish a *Report on AI in Financial Services*, the CMA published a report on *Algorithms: How they can reduce competition and harm consumers*, Ofcom commissioned a report on the *Use of AI in Online Content Moderation*, and the MHRA has launched the Software and AI as a Medical Device Change Programme.

[5] For example, the UK Government's Office for AI published *Guidelines for AI Procurement* and *Guidance on Ethics, Transparency and Accountability Framework for Automated Decision-Making* in the public sector, and the UK Ministry of Defence (MoD) published a policy statement on *Ambitious, safe, responsible: our approach to the delivery of AI-enabled capability in Defence* which is relevant to MoD suppliers.

Government is exploring options for encouraging regulatory coordination through platforms such as the Digital Regulation Cooperation Forum (DRCF)[6] to ensure coherence among the regulators and to support innovation.

The UK Government recognises that regulators will need access to the necessary skills and expertise to effectively regulate AI. Although we have seen a number of regulators investing in their AI capabilities over the past several months, it is currently not clear that the regulators will be able to keep the pace with investment in AI capabilities from the business sector. The Policy Paper mentioned that the UK Government will explore the possibility of pooling resources and capabilities among multiple regulators, as well as the options for secondments from businesses and academia, to help regulators access the skills and expertise needed.

Comparison to the European Commission's AI Act Proposal

The European Commission's proposal for an AI Act published in April 2021 and the UK Government's Policy Paper set out differing views for regulation of AI in Europe and show one of the first major divergences in regulatory approach between the EU and the UK post-Brexit. We have summarised some of the major differences:

- **Sector-specific approach**: Unlike the EU's AI Act proposal, the Policy Paper sets out a de-centralised approach to AI regulation. The UK Government rejected the idea of creating a single regulator with a new mandate and enforcement powers responsible for regulating AI across all sectors. Instead, the UK Government plans to leverage the experience and expertise of existing regulators and ask them to issue guidance to highlight the relevant regulatory requirements applicable to businesses they regulate (such as any requirements for meeting sector-specific licences or standards, or appointing named individuals to assume particular responsibilities). The UK Government also hopes that this de-centralised approach will be more adaptable to technological change.
- **No central list of prohibited or high-risk use cases**: The EU's AI Act proposal includes a list of prohibited AI practices that are unacceptable in all circumstances (including certain uses of real-time remote biometric identification) as well as a list of high-risk AI systems which have to undergo a conformity assessment and comply with strict requirements in the AI Act. On the other hand, the Policy Paper does not seek to ban specific uses of AI but will leave it up to regulators to decide if the use of AI in specific scenario should not be allowed or should be subject to higher regulatory burden.
- **No new legislation (at least for now)**: The EU's AI Act is a proposal for a new regulation which would be directly applicable in all EU Member States. On the other hand, the UK Government proposes to initially put the cross-sectoral principles on a non-statutory footing, for example, by issuing executive guidance or specific mandate to regulators without introducing new legislation. However, the UK Government has not ruled out proposing new legislation where and when needed to ensure effectiveness of the new

[6] The DRCF formed in July 2020 and comprises the CMA, ICO, Ofcom, and since April 2021 also the FCA.

regulatory framework. Alongside the Policy Paper, the UK Government is proposing changes to existing UK legislation to make the use of AI in the UK easier (such as proposing amendments to Article 22 of the UK General Data Protection Regulation[7] or introducing a new text and data mining exemption for any purpose in the Copyright, Designs and Patents Act 1988[8]).

Next steps

The UK Government is seeking initial views from stakeholders on the proposal set out in the Policy Paper. The call for views is open until 26 September 2022.

Following the call for views, the UK Government is expected to publish an AI White Paper in late 2022 which will set out more concrete proposals for AI regulation in the UK.

WHAT SHOULD BUSINESSES DO NOW?

1. Senior leaders should consider who is responsible for AI governance and risk management strategy within their organisation.
2. Businesses should review their internal AI strategy and the proposed principles and consider what steps they will need to take to align the strategy with the new AI regulatory frameworks that are emerging in the EU, UK and elsewhere[9].
3. Organisations that use AI in the UK or licence AI for use in the UK should monitor the upcoming AI White Paper and announcements from the relevant regulators about how they will interpret, implement and enforce the cross-sectoral principles.

[7] The UK Government proposed amendments to the UK General Data Protection Regulation in the Data Protection and Digital Information Bill introduced in the UK Parliament on 18 July 2022.

[8] *Artificial Intelligence and Intellectual Property: copyright and patents: Government response to consultation*.

[9] For example, the US Office of the Comptroller of the Currency (OCC) outlined some of the supervisory expectations for how the banks it regulates should manage risks associated with AI.