A Glimpse Into The Potential Future Of AI Regulation

By **Byungkwon Lim, Gary Murphy, Friedrich Popp and James Amler** (May 14, 2019)

Byungkwon Lim

Gary Murphy

Friedrich Popp

James Amler

On April 10, 2019, Democrats in Congress introduced legislation, referred to as the Algorithmic Accountability Act of 2019, seeking to enhance oversight of artificial intelligence and data privacy.[1] If passed, the AAA would direct the Federal Trade Commission to "require entities that use, store, or share personal information" to conduct impact assessments of any automated decision systems or any repositories of consumer information that are deemed to be "high risk." In this way, the AAA creates a regulatory framework designed to reduce the risk that (1) AI systems lead to inaccurate, unfair, biased, or discriminatory outcomes for consumers, or (2) personal data of consumers is improperly disseminated.

These protections proposed in the AAA in some respects mirror protections already enacted in Europe under the General Data Protection Regulation. Although the AAA faces uncertain prospects of becoming law given current political realities, it offers a glimpse into potential future U.S. regulation of AI on the federal or state level. Unfortunately, the view from the near future appears to be one of uncertainty, both in terms of who will regulate and what exactly is required.

Summary of the Proposed AAA Bill

At its core, the AAA is a direction to the FTC to implement more stringent oversight of how personal data is stored, used and analyzed by businesses and other entities that have access to large volumes of consumer data. The AAA does not expressly provide for a private right of action, nor extraterritorial jurisdiction, but rather is designed to be enforced "in the same manner, by the same means, and with the same jurisdiction, powers, and duties as" other FTC statutory provisions.[2]

If the AAA becomes law, within two years of its enactment, the FTC would be obligated to implement regulations requiring covered entities to conduct periodic "impact assessments" of existing consumer AI systems or data systems deemed "high risk" either for data breaches or for the use of algorithmic determinations that are potentially unfair or biased. New consumer AI systems or information systems deemed "high risk" would require an impact assessment prior to implementation.

The AAA specifies that the impact assessments should be conducted, when possible, "in consultation with external third parties, including independent auditors and independent technology experts." Once it has completed the impact assessment of either the AI system or the information system, the business would be required to "reasonably address in a timely manner the results of the impact assessments."

Entities covered by the AAA include businesses over which the FTC has jurisdiction and that meet any of the following criteria:[3]

- had greater than $50 million in average gross revenue over the past three years;

- possesses or controls data for over one million consumers or consumer devices; or

- is a data broker or other commercial entity that sells, trades, or provides third-party access to personal information.

Oversight of Automated Decision Systems

The AAA initially casts a wide net of activity within its scope. It defines "automated decision system" broadly to mean "a computational process, including one derived from machine learning, statistics, or other data processing or artificial intelligence techniques, that makes a decision or facilitates human decision making, that impacts consumers."

This definition would include a broad range of systems used to process consumer information, to provide services to consumers, or in any other way that "impacts" consumers. The bill appears to target just about any automated computational technology, not only ones in which artificial intelligence actually makes a determination, but even so-called "augmented intelligence," where a computer process serves to assist and enhance human determinations.

Only those automated decision systems deemed to be "high risk" would be subject to the required impact assessment. The AAA defines ''high-risk automated decision system'' as one that warrants enhanced scrutiny, either because the types of data inputs are particularly sensitive or because the outputs of the system are especially impactful on consumers. An automated decision system could be designated high risk based upon any of the following criteria:

> (A) taking into account the novelty of the technology used and the nature, scope, context, and purpose of the automated decision system, [it] poses a significant risk (i) to the privacy or security of personal information of consumers; or (ii) of resulting in or contributing to inaccurate, unfair, biased, or discriminatory decisions impacting consumers;
>
> (B) [it] makes decisions, or facilitates human decision making, based on systematic and extensive evaluations of consumers, including attempts to analyze or predict sensitive aspects of their lives, such as their work performance, economic situation, health, personal preferences, interests, behavior, location, or movements, that—(i) alter legal rights of consumers; or (ii) otherwise significantly impact consumers;
>
> (C) [it] involves the personal information of a significant number of consumers regarding race, color, national origin, political opinions, religion, trade union membership, genetic data, biometric data, health, gender, gender identity, sexuality, sexual orientation, criminal convictions, or arrests;
>
> (D) [it] systematically monitors a large, publicly accessible physical place; or

(E) [it] meets any other criteria established by the [FTC] in regulations issued under section 3(b)(1) [of the AAA].

Any such high-risk automated decision system would require an impact assessment to evaluate the system's development process, design and training data for "impacts on accuracy, fairness, bias, discrimination, privacy, and security." The AAA specifies that the impact assessment must include:

- a detailed description of the system's functionality, design, training data, and purpose;

- a cost/benefit assessment of the system in light of its purpose, its data storage and minimization practices, its transparency to consumers, and opportunities for consumers to correct or object to its results;

- an evaluation of the data privacy risks;

- an evaluation of "the risks that the automated decision system may result in or contribute to inaccurate, unfair, biased, or discriminatory decisions impacting consumers"; and

- a plan for technological and physical safeguards to minimize the risks to data privacy or harmful automated results.

Oversight of Information Systems

Under the AAA, a "data protection impact assessment" is required to evaluate whether a high-risk information system protects the privacy and security of personal information stored in or processed by the system.

The AAA defines "information system" extremely broadly to include any process (other than an automated decision system) "that involves personal information" including "collection, recording, organization, structuring, storage, alteration, retrieval, consultation, use, sharing, disclosure, dissemination, combination, restriction, erasure, or destruction of personal information." But the required assessments under the AAA only apply to information systems designated as "high-risk."

The criteria by which an information system is deemed "high-risk" are somewhat nebulous, but are set forth in three categories: the technology is novel or has insufficient data security; the consumer data involves sensitive personal information such as race, religion, political opinions, health data, etc.; or the data involves systematic monitoring of a large public space. The AAA also permits the FTC to establish additional criteria for deeming an

information system as high-risk.

Challenges for Compliance With the AAA

Businesses would face significant uncertainty in seeking to comply with the AAA's proposed requirements for implementing AI systems and utilizing consumer data. The AAA only calls for impact assessments of "high-risk" AI systems and information systems, which is a potentially helpful distinction. However, the categories that define high-risk are open ended and will be difficult to delineate. Moreover, the FTC would enjoy discretion to create new criteria for identifying high-risk AI or information systems.

Another major challenge for companies will be to provide the FTC with sufficient documentation and description of their AI systems including how they work and how they train personnel to use the system. A persistent operational issue encountered with AI and machine learning is that developers often cannot reliably predict the outcome in advance nor effectively explain the results ex-post. An algorithm might be very effective for its purpose, but the more opaque its functioning, the more likely it will be viewed by a regulator with skepticism — or even suspicion. There is already precedent for enforcement of anti-discrimination laws on the basis of outcomes alone: the Consumer Financial Protection Bureau's actions against auto lenders for discriminatory pricing on race.[4]

Comparison to GDPR

The European General Data Protection Regulation took effect in 2018 and applies to the processing of personal data of individuals in the European Union irrespective of their citizenship.[5] For this purpose, personal data includes any information relating to an identified or identifiable natural person.

The GDPR also establishes certain special categories of personal data that benefit from specific protection. These categories include data on matters such as racial or ethnic origin, political opinions, religious or philosophical beliefs, trade union membership, genetic and biometric information, health information, sexuality, sexual orientation, criminal offenses, and criminal convictions.

The GDPR applies to businesses established in the European Union as well as controllers and processors of personal data established outside the Union if they sell goods or provide services to individuals in the EU or monitor the behavior of individuals in the EU.

The AAA is similar to the GDPR in a number of respects. For example, the GDPR requires that in scenarios where the processing is likely to result in a high risk for the fundamental right to data protection or the prohibition of discrimination, prior to such processing every controller of personal data must conduct a documented data protection impact assessment.

A DPIA is required in particular for personal data collected and used in the following contexts: (1) a systematic and extensive evaluation relating to natural persons that is based on automated processing and on which decisions are based that produce legal effects concerning the natural person or that similarly significantly affect the natural person; (2) processing on a large scale of identified special categories of data; and (3) a systematic monitoring of a publicly accessible area on a large scale. Supervisory authorities guide in nonexhaustive lists whether a DPIA may be required for a certain type of processing. The periodic assessment is also required to address a number of factors that, while not identical to those set forth for impact assessments under the AAA, are fairly similar in scope and purpose.

If the assessment indicates that the processing would result in a high risk, the GDPR requires the controller to consult with the supervisory authority. The authority determines inter alia on the basis of the DPIA whether the controller sufficiently identified or mitigated the risk and advises within a short time frame whether the envisaged processing would infringe the law.

Certain aspects of the GDPR are not addressed explicitly in the text of the AAA but may provide a sense of the regulations that could be adopted by the FTC if the AAA were passed. For example, the GDPR grants enhanced protections and rights to data subjects in the case of profiling and automated decision-making. These include specific requirements regarding transparency and fairness, greater accountability, and legal bases for processing, as well as enhanced rights for data subjects (including the right to oppose profiling or the right to not become subject of a significant decision that is based solely on automated processing).

Is the AAA a Glimpse of the Future?

Eventually, the AAA itself could be enacted, for example, if a Democratic administration comes to power in 2021 with congressional majorities. Yet even if the AAA does not pass Congress, the proposed bill could provide a framework for executive action by a Democratic president to direct the FTC to regulate AI and information systems. Also, it is possible that some high-profile AI failure or data breach could generate a groundswell of bipartisan support for federal action.

If action fails at the federal level, state-based legislation could gain momentum, and states could adopt their own versions of the AAA. For example, in January 2019 a bill was introduced in Washington state that follows the GDPR model of prohibiting algorithmic-based decisions from discriminating on the basis of any protected classifications.[6] The Washington bill focuses more upon regulating the use of automated decision systems by state agencies and lacks the AAA's specific requirements for auditing and reporting directed at large private businesses.

Despite the introduction of the AAA in Congress, various state privacy protection laws may ultimately lead the way in implementing restrictions on profiling and algorithmic discrimination. Many in the industry view it as inevitable that many, if not all, states will adopt some sort of privacy protection laws. The expectation is that states would strongly object to any attempt for federal preemption on this, as states have been primarily responsible for consumer protection for their residents. In fact the AAA itself expressly clarifies that it does not preempt any state law. This could present a nightmare scenario for businesses, potentially needing to comply with a federal law in addition to 50 state laws (plus D.C. and Puerto Rico).

Conclusion

Although the AAA faces uncertain prospects of becoming law given the current political realities, it offers a glimpse into potential future regulation of AI. The AAA creates a regulatory framework — similar to the GDPR — designed to protect consumers from AI systems that are inaccurate, unfair, biased or discriminatory, and from improper use of consumer personal data. However, even if the AAA were to be passed and/or used to set a uniform standard, businesses would face significant uncertainty in implementation as well as likely conflicting regulation in other jurisdictions.

Byungkwon Lim is a partner, Gary Murphy is counsel, Friedrich Popp is a senior associate and James Amler is counsel at Debevoise & Plimpton LLP.

[1] Algorithmic Accountability Act of 2019, S. 1108, 116th Cong. (2019). The bill was proposed in the U.S. Senate by Sen. Ron Wyden (D-OR) and Sen. Corey Booker (D-NJ). Id. The bill was sponsored in the U.S. House of Representatives by Yvette Clarke (D-NY). Algorithmic Accountability Act of 2019, H.R. 2231, 116th Cong. (2019).

[2] The FTC has jurisdiction to enforce a prohibition over "unfair or deceptive acts or practices" in the Federal Trade Commission Act, 15 U.S.C. § 45(a)(1) (2006), as well as cases involving foreign commerce or conduct that implicate federal antitrust laws. See U.S. Dep't of Justice & Federal Trade Comm'n, Antitrust Guidelines for International Enforcement and Cooperation (2017), https://www.ftc.gov/system/files/documents/public_statements/1049863/international_guidelines_2017.pdf.

[3] The FTC's jurisdiction is defined under section 5(a)(2) of the Federal Trade Commission Act. 15 U.S.C. § 45(a)(2). The FTC has broad statutory jurisdiction over all persons or entities using unfair methods of competition in or affecting commerce, with exceptions for some highly regulated sectors including banks, common carriers, and air carriers. See id.

[4] Press Release, Consumer Fin. Prot. Bureau, CFPB to Hold Auto Lenders Accountable for Illegal Discriminatory Markup: Bureau Provides Guidance on Fair Lending Practices to Indirect Auto Lenders (Mar. 21, 2013), https://www.consumerfinance.gov/about-us/newsroom/consumer-financial-protection-bureau-to-hold-auto-lenders-accountable-for-illegal-discriminatory-markup/. On May 21, 2018, the President signed a joint resolution passed by Congress disapproving the CFPB Bulletin specifying disparate impact as a theory for auto lender liability. Id.

[5] Commission Regulation 2016/679, 2016 O.J. (L 119).

[6] S. 5527, 66th Leg., Reg. Sess. (Wash. 2019).