



February 26, 2025

Re: SB 936: An Act Concerning Artificial Intelligence - Favorable with Amendments

Dear Honorable Members of Senate Finance Committee,

Consumer Reports¹ writes to you regarding SB 936, which we support with amendments. This bill is focused on a critically important issue—ensuring that artificial intelligence systems used in high-stakes decisions about Marylanders are transparent, free of bias, and that residents retain some autonomy and recourse.

If the committee adopts the amendments presented today, this bill could act as an essential patch for Maryland’s existing civil rights and consumer protection laws in the A.I. era. While these laws clearly apply to A.I. products, the “black box” nature of these systems and their ability to partially stand in for human decision-makers with intent can make cases difficult to bring in practice.

This bill, if amended, provides a measured solution to that problem. It would apply to companies developing and using AI systems to help make critical decisions about Marylanders—such as who gets a job, who gets into their dream college, who is selected for an apartment, who gets access to which medical treatments, and more. SB 936 would require companies to adopt industry best practices for assessing the risks associated with their products; would require developers to share information with the companies that buy their products; and would require companies to provide information to consumers so that they can understand how AI systems are making high-stakes decisions about them, and can vindicate their rights under other laws.

Predictive AI systems can be biased, error-prone, or snake oil

When companies use predictive artificial intelligence to make critical decisions about Marylanders—such as whether a consumer gets a job, is offered their dream apartment, or qualifies for certain health services—consumers may be subject to bias, or erroneous conclusions. For example:

¹ Founded in 1936, Consumer Reports (CR) is an independent, nonprofit and nonpartisan organization that works with consumers to create a fair and just marketplace. Known for its rigorous testing and ratings of products, CR advocates for laws and company practices that put consumers first. CR is dedicated to amplifying the voices of consumers to promote safety, digital rights, financial fairness, and sustainability. The organization surveys millions of Americans every year, reports extensively on the challenges and opportunities for today's consumers, and provides ad-free content and tools to 6 million members across the U.S.

- One resume screening program identified two factors as the best predictors of future job performance: having played high school lacrosse and being named Jared.²
- Another employment assessment provided high scores in English proficiency even when questions were answered exclusively in German.³
- A health care algorithm used widely by hospitals to identify which patients would receive additional care was found by independent researchers to be biased against Black patients; in attempting to predict which patients would become the sickest, it instead predicted who would spend the most money care.⁴
- A sepsis-prediction algorithm used by many hospitals nationally was found to not be nearly as accurate as the company selling it had claimed—and only slightly more accurate than simply flipping a coin.⁵

Unfortunately, it's all too rare for these issues to come to light. When other products malfunction—such as when a toaster explodes, or airbags in a car fail to deploy—it is immediately apparent to consumers and enforcers that there's a problem. When a predictive artificial intelligence product fails, the developer can, in most cases, keep that quiet. No other entity or person has good insight into the product they've developed. The deployer using the AI system may not really know how the system works, or that the system is faulty. The consumer whose life is ultimately impacted also doesn't know if the system is discriminatory or error-prone—and may not know that an AI system is being used at all.

SB 936 as written has serious weaknesses; amendments greatly improve it

As written, SB 936 has loopholes that could allow many companies to side-step accountability entirely. Here, we will expand on just a few.

“High risk artificial intelligence system” is currently defined as “an artificial intelligence system that is specifically intended to autonomously make, or be a substantial factor in making a consequential decision.” This narrow definition, focused on the developer’s “specific intent,” would allow developers to sidestep this law entirely if they simply market their tool as “assisting” in decisions, or add “only to be used with human supervision” to their documentation.

Additionally, the bill’s current definition of “substantial factor” is overly complex, and generates the need for yet another term and definition: “principal basis.” The way these two definitions

² Dave Gershgorn, Quartz, “Companies are on the hook if their hiring algorithms are biased,” (Oct 22, 2018) <https://qz.com/1427621/companies-are-on-the-hook-if-their-hiring-algorithms-are-biased>

³ Sheridan Wall, Hilke Schellmann, MIT Technology Review, “We tested AI interview tools. Here’s what we found,” (July 7, 2021) <https://www.technologyreview.com/2021/07/07/1027916/we-tested-ai-interview-tools/>

⁴ Ziad Obermeyer, Brian Powers, Christine Vogeli, Sendhil Mullainathan, *Science*, “Dissecting racial bias in an algorithm used to manage the health of populations” <https://www.science.org/doi/10.1126/science.aax2342>

⁵ Arvind Narayanan and Sayash Kapoor, Financial Review, (Sept. 13, 2024) <https://www.afr.com/technology/snake-oil-don-t-believe-the-artificial-intelligence-hype-20240909-p5k93y>

interlock with each other and with the definition of “high risk artificial intelligence system” means that, in practice, companies could simply assign humans to rubber stamp algorithmic recommendations in order to be released from the law entirely. Having humans approve AI-generated recommendations is not a meaningful alternative to the provisions in this bill, in part because ample research suggests that humans tend to view automated systems as authoritative and trustworthy, and are inclined to defer to their recommendations—even when they suspect the system is malfunctioning.⁶

The bill as written also releases companies from nearly all of its requirements if they are in “conformity” with the latest version of the AI Risk Management Framework published by the National Institute of Standards and Technology. This seriously misunderstands the purpose of the AI Risk Management Framework, which is not intended to and cannot take the place of regulation.

Overview of selected amendments

Senator Hester today presents several thoughtful amendments that would streamline the bill and ensure it doesn’t contain unnecessary loopholes that would undercut the intent of the law. Here, we will touch several of those amendments and why they are important.

Removing loopholes and unnecessary, vague exemptions to the definition of “high risk artificial intelligence:” As written, most provisions of this law are scoped by the definition of “high-risk artificial intelligence system” and a few other key definitions. As mentioned above the “specifically intended” language in this definition as currently written creates an easy way for companies selling predictive AI products to side step most provisions of this bill entirely.

This definition also contains a couple unnecessary exemptions that are ripe for abuse. For example, “narrow procedural tasks” are exempted, but that term is undefined, leaving ambiguous whether core activities this law should cover—such as screening and scoring resumes, or housing applicants—could be considered “narrow procedural tasks.” The exemption is also unnecessary, since the law already enumerates and exempts elsewhere the types of technologies that execute narrow procedural tasks, such as spell-check, spreadsheets, databases, and more. The “perform a preparatory task to an assessment relevant to a consequential decision,” exemption is also ripe for abuse. To use the same examples above, AI screening and scoring of job or housing applicants could also be considered a preparatory task to the ultimate “assessment,” made by a human who is considering a couple other factors in his or her ultimate decision. Nevertheless, these “preparatory tasks” could be highly influential on the outcome of the decision.

⁶ See, e.g., Danielle Keats Citron, Washington University Law Review, ‘Technological Due Process,’ 2008 at 1271–72; https://openscholarship.wustl.edu/cgi/viewcontent.cgi?article=1166&context=law_lawreview

The amendments deal with both of these issues by changing the definition of “high risk artificial intelligence” to read:

(1) “HIGH-RISK ARTIFICIAL INTELLIGENCE SYSTEM” MEANS AN ARTIFICIAL INTELLIGENCE SYSTEM THAT, WHEN DEPLOYED, MAKES, OR IS A SUBSTANTIAL FACTOR IN MAKING, A CONSEQUENTIAL DECISION.

And by cutting exemptions (2)(I)(1-4).

Adds “cost or terms” to the definition of “consequential decision”: The cost of housing, insurance, lending, and the material terms of employment, housing, lending, and more, can be functionally just as important to a consumer’s ability to access these opportunities as provision or denial. To account for that, “cost or terms” has been added to the definition of “consequential decision.” We would note, however, that this definition is still scoped to the most important decisions that have “a materially legal or similarly significant effect.”

Removes “duty of care” framing for algorithmic discrimination: As written, this bill addresses algorithmic discrimination by requiring that developers and deployers “use reasonable care to protect consumers from known and reasonably foreseeable risks” of algorithmic discrimination.

This framework is problematic for a few reasons. First, it suggests that “algorithmic discrimination” is somehow less harmful or less important than other forms of discrimination, which state and federal laws prohibit. Second, it risks confusion; if a company uses a discriminatory AI hiring tool and it’s brought to the attention of an enforcer, they may be sued under both a new algorithmic discrimination chapter, and existing employment discrimination laws. The company may be in compliance with the algorithmic discrimination statute, and in violation of employment antidiscrimination law. This dynamic will make such discrimination cases more complex to litigate, more confusing to juries, and may muddy the waters on what would have once been a straightforward finding of employment discrimination. Lastly, there’s a concerning possibility that court decisions under this “duty of care” approach could bleed into interpretation of existing civil rights statutes, weakening them.

For those reasons, amendments eliminate the “duty of care” language.

Narrow the overbroad exemption to consumer’s right to appeal: As written, companies do not have to provide a right to appeal when doing so “is not in the best interest of the consumer.” This is an overbroad exemption that will result in consumers being unfairly denied the right to appeal. Furthermore, consumers are better situated than companies to evaluate whether exercising their right to appeal is in their own best interest. When an appeals process is not in a consumer’s best interest, they can always choose not to pursue it. The amendments narrow this exemption to situations where offering the right to appeal would pose a risk to the life or safety of the consumer.

Removes NIST RMF “conformity” exemption: As written, a company would be deemed to be in compliance with the bill if it is in “conformity” with the AI Risk Management Framework published by the National Institute of Standards and Technology (NIST RMF). This is an odd provision, and would seem to misunderstand the purpose of the NIST RMF. The NIST RMF has many laudable elements but it was not written as a stand in for regulation; it was not written with compliance in mind. Furthermore, the text of the RMF itself says it does not prescribe risk tolerance. Indeed, it reads:

While the AI RMF can be used to prioritize risk, it does not prescribe risk tolerance. Risk tolerance refers to the organization’s or AI actor’s readiness to bear the risk in order to achieve its objectives. Risk tolerance can be influenced by legal or regulatory requirements (Adapted from: ISO GUIDE 73). Risk tolerance and the level of risk that is acceptable to organizations or society are highly contextual and application and use-case specific. Risk tolerances can be influenced by policies and norms established by AI system owners, organizations, industries, communities, or policy makers.⁷

A company could decide it has a very high tolerance for the risk of bias and still be in “conformity” with the NIST RMF. The amendments remove the provisions that allow companies to opt-out of all the provisions in the bill—including all of the disclosure provisions—by aligning with the NIST RMF.

With the committee amendments presented today, many issues have been resolved. We urge the committee to refer the bill favorably with amendments. If passed into law, this bill will become a critical tool for consumers to protect themselves in the AI era.

We thank Senator Hester for her leadership on this issue, and we thank the committee for its consideration.

Sincerely,
Grace Gedye
Policy Analyst, Consumer Reports

⁷ ‘Artificial Intelligence Risk Management Framework (AI RMF 1.0)’, *National Institute of Standards and Technology, U.S. Department of Commerce*, January 2023, accessed at: <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>

