February 11, 2025

Re: House Bill 2094: High-risk artificial intelligence; development, deployment, and use, civil penalties – oppose unless amended

Dear Honorable Senators of the General Laws and Technology Committee,

Consumer Reports[1] writes in respectful opposition to H.B. 2904 unless the bill is amended. The bill is focused on an important issue—ensuring that artificial intelligence systems used in high-stakes decisions are transparent, free of bias, and that Virginians retain some autonomy and recourse. However, we are concerned that significant loopholes will undermine this goal, leading to low compliance and leaving Virginians in the dark. As such, the bill should be substantially strengthened before it is enacted and should not be approved in its current form.

When companies use predictive artificial intelligence and automated decision systems to make critical decisions about Virginians—such as whether a resident gets a job, is offered their dream apartment, qualifies for certain health services, or lands a spot in their top-choice college—consumers may be subject to bias, or erroneous inferences and conclusions.

Investigations by journalists, researchers, and the testimony of experts helping companies evaluate predictive AI products have produced worrying results. 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 assessment provided high scores in English proficiency even when questions were answered exclusively in German.[3] A sepsis-prediction algorithm sold widely to hospitals across the country was found, when evaluated by independent

[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.

[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/

researchers, to not be nearly as accurate as the company selling it had claimed and only slightly more accurate than simply flipping a coin.[4]

Our research suggests that consumers are uncomfortable with the use of AI in the types of high-stakes decisions that this bill covers. In May of 2024, Consumer Reports conducted a nationally representative study of 2,022 U.S. adults focused on the use of AI and algorithms in consequential decisions.[5] The survey was administered by NORC at the University of Chicago. When asked how they feel about the use of AI and algorithms in a variety of situations—such as banks using algorithms to make underwriting decisions, landlords using AI to screen potential tenants, hospitals using AI to help make diagnoses—a majority of Americans said they were uncomfortable with each scenario.

Ensuring that companies making and using predictive AI and automated decision systems do basic due diligence is critical; informing consumers about how decisions are being made about them with their data is essential to their ability to exercise their existing rights under civil rights laws, consumer protection laws, and more. As written, however, we are concerned that H.B. 2094 contains loopholes, exemptions, and shields for companies that would lead to high levels of non-compliance.

We therefore have the following suggestions.

Fix loophole in definition of 'high risk AI' and 'substantial factor' so that companies cannot escape responsibility: Currently, most provisions in the law only apply to developers and deployers of "high-risk artificial intelligence systems" defined in part as a 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. This definition is far narrower than the definition under a similar, existing law passed by Colorado.

The definition of "substantial factor" would allow companies making high-stakes decisions about Virginians to evade accountability, so long as they can justify to themselves that they don't use it as the "principal basis" for a decision. Once companies decide they are not using a high-risk AI tool as a "principal basis" for their consequential decisions, they are also exempt from the law's disclosure requirements, making it exceedingly difficult for any enforcer in practice to identify and challenge a company's erroneous interpretation of the term.

---

[4] 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

[5] Consumer Reports Survey Group, A.I./Algorithmic Decision-making: Consumer Reports Nationally Representative Phone and Internet Survey, (July 9th, 2024)
https://advocacy.consumerreports.org/wp-content/uploads/2024/07/CR-AES-AI-Algorithms-Report-7.25.24.pdf

Similar language was used to define "substantial factor" in New York City's Local Law 144. Researchers at Cornell University assessed compliance with the law and found extremely low levels of affirmative compliance, which they attributed in part to the regulations' definition of "substantial factor."[6]

Lastly, we suggest that exemptions (i)-(iv) be cut from the definition of high risk AI. Any one of these exemptions could undermine the law significantly. For example, exemption (i), "narrow procedural task" 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 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.

We suggest:

- Amending the bill's definition of high-risk artificial intelligence system to read: "Any artificial intelligence system that, when deployed, makes, or is a substantial factor in making, a consequential decision."
- Amending the bill's definition of high-risk artificial intelligence system to read: "Substantial factor means a factor that: (i) assists in making a consequential decision; (ii) is capable of altering the outcome of a consequential decision; (iii) is generated by an artificial intelligence system."
- Cutting exemptions (i)-(iv) in the definition of "High risk artificial intelligence system"

Require developers to test for disparate impact; prohibit developers from selling products that produce unjustified disparate impact: Currently, the only responsibilities the bill places on developers of high risk AI technology are documentation requirements. Developers of high risk AI systems should be required to test their products for disparate impact before putting them on the market, and should not put their products on the market if the testing reveals unjustified disparate impact. We suggest adding this requirement.

***Strengthen antidiscrimination protections*:** While this bill purports to address discrimination, it makes a significant departure from landmark anti-discrimination laws which prohibit discrimination outright. This bill instead adopts a "duty of care" approach, requiring that companies developing and using high-risk AI "use a reasonable duty of care to protect consumers from any known or reasonably foreseeable risks of algorithmic discrimination." It also provides companies with a rebuttable presumption shielding them from liability; if

[6] Wright, L., Muenster, R. M., Vecchione, B., Qu, T., Cai, P., Smith, A., … & Matias, J. N. (2024). Null Compliance: NYC Local Law 144 and the challenges of algorithm accountability. In The 2024 ACM Conference on Fairness, Accountability, and Transparency.

companies fulfill the largely procedural and documentation requirements of the bill, they are presumed to have met their duty of care—regardless of whether discrimination actually did occur.

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 this 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, we are concerned that court decisions under this algorithmic discrimination bill could bleed into interpretation of existing civil rights statutes, weakening them.

We suggest removing the duty of care, and instead adding a prohibition against deployers using high risk AI tools in such a manner that causes discrimination, or developers selling or placing into the stream of commerce products that discriminate. Short of that, we would suggest eliminating the rebuttable presumption, which would allow the duty of care to remain flexible over time, as testing and debiasing high risk AI products becomes increasingly simple and cheap and what constitutes a "reasonable duty of care" evolves.

Remove overbroad exemptions to what companies must disclose: Many of the substantive requirements of this bill hinge on companies disclosing information to one another, to consumers, and to the Attorney General. Currently the bill permits companies to withhold trade secrets and also "other confidential or proprietary information." If companies can unilaterally decide to withhold important information, it will hamper the efficacy of this law.

It is concerning companies are permitted to withhold trade secrets from the Attorney General; the Attorney General's office is not a competitive threat to companies, and as the sole enforcer of this law it is critical that that office is able to gather the information necessary to investigate wrongdoing.

We suggest:

- Cutting "other confidential or proprietary information" from disclosure exemptions, as these are ill defined categories that go well beyond existing trade secret exemptions and will be abused.
- Eliminating companies' ability to withhold trade secrets when providing information directly to the Attorney General

Clarify consumer disclosures and remove exemption from consumers' right to appeal: This bill requires companies to disclose some information to consumers, including when a company is about to use a high risk AI system to make a consequential decision about a Virginian (also referred to as a "pre-use notice"), and an explanation after a high risk AI system has made a consequential decision about a Virginian. So long as other loopholes in the law are addressed, these disclosures will provide consumers in Virginia with meaningful transparency, and will enable them to make use of other rights, such as the right to appeal.

Currently, the law requires deployers to provide consumers with a pre-use notice when a "deployer uses a high-risk artificial intelligence system to interact with a consumer." Consumers often do not directly interact with high-risk AI. For example, when a landlord uses high-risk AI to rate a housing applicant, the applicant typically interacts with the landlord directly and does not interface with the AI system. We suggest modifying this language to clarify that the pre-use notice should be delivered when a deployer "deploys a high-risk artificial intelligence system to make, or be a substantial factor in making a consequential decision concerning a consumer." This language would also bring this bill into alignment with a similar existing law in Colorado.

Another important provision in this bill is consumers' right to appeal. Currently, 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.

Reconsider overbroad exemptions and remove cure provision: We caution the committee to carefully consider the broad exemptions for entire industries. It is not clear to us why, for example, companies covered by HIPAA—which serves a different purpose than this bill—should be exempt. It is also unclear to us why high-risk AI systems "acquired" by federal agencies should be exempt from the bill.

Lastly, we would suggest eliminating the cure provision. Attorneys General already have the discretion to no longer pursue a case if the facts are no longer compelling. Cure provisions create perverse incentives for compliance and hamstring regulators' ability to hold wrongdoers accountable.

We appreciate the intent of this bill and Delegate. Maldonado's leadership on the issue; it is an important topic to tackle. However, given the current draft of the bill, we are respectfully opposed, unless the bill is amended.

Sincerely,
Grace Gedye
Policy Analyst, Consumer Reports