Testimony of Matthew Scherer, Senior Policy Counsel for Workers' Rights & Technology, Center for Democracy & Technology
Before the Maryland House of Delegates Economic Matters Committee
March 4, 2025

About CDT

The Center for Democracy & Technology (CDT) is a nonprofit, nonpartisan organization fighting to advance civil rights and civil liberties in the digital age. CDT works on many issues touching on various aspects of artificial intelligence, algorithmic systems, and related technologies. It also has workstreams specifically focused on digital rights in specific fields relevant to automated decisions, including workers' rights and disability rights.

Introduction

We want to thank Delegate Qi for this thoughtful bill that would provide crucial protections to consumers and workers whose lives are increasingly driven by hidden algorithmic and AI decision systems. HB 1331 would provide some crucial transparency requirements and guardrails on the use of AI, benefiting consumers subjected to such decisions, while avoiding many of the loopholes and other pitfalls contained in similar bills pending in other states. That said, a few amendments are needed to ensure the bill provides adequate protection for Marylanders.

My testimony begins with an overview of the problem of hidden algorithmic decisions, then rebuts several arguments that we know you and your colleagues in the legislature are hearing and will continue to hear from industry groups. The last section of the testimony lays out changes that are needed to ensure the bill achieves its goals.

Automated decision systems are frequently hidden from the workers and consumers they affect and carry significant risks, particularly to those from vulnerable and marginalized communities

Artificial intelligence (AI) has seen rapid and genuine progress in recent years. Generative AI systems have allowed for the introduction of consumer-facing chatbots that perform leaps and bounds better than their predecessors from just a few years ago. AI is also showing promise in the areas of scientific and medical research. States can, and should, encourage innovation in those and other areas of AI research that clearly advance the public interest.

But not all applications of AI are beneficial, particularly when it comes to their impacts on vulnerable workers and consumers. Companies increasingly use AI-powered ADSs when determining who to hire, mortgage rates for bank customers, who can access public benefits like SNAP or Medicaid, and how much we all pay for life's necessities like health care and rent.

Unfortunately, there is considerable evidence that many such ADSs are biased (or simply do not work) and that removing such bias is quite difficult. At the same time, the lack of transparency surrounding companies' ADS practices means that workers, consumers, and regulators only rarely catch glimpses of when, how, and on whom companies use these tools.

We do not have a clear picture of how, nor how many, companies use these technologies because information disclosure about ADS use by companies is sporadic and inconsistent, if there is disclosure of the ADS use at all. That said, considerable anecdotal evidence suggests that the practice is already widespread. Surveys of companies indicate that anywhere from one-third to the vast majority already use ADSs in recruitment and hiring alone. But we often don't know *which* companies are using these tools, nor which workers and consumers are being affected by them.

Stories about harmful uses of ADSs have come to light thanks to whistleblowers and investigative journalism. ProPublica has published a trio of reports on how the health care giant Cigna secretly used an algorithm to mass-reject policyholders' claims—and then threatened to fire a physician who pushed back. But consumers and workers should not be forced to rely on whistleblowers and nonprofit news outlets to bring these issues to light, nor to fight harmful ADSs.

The significant information asymmetry surrounding ADSs provides a strong reason for regulation because existing civil rights, labor, and consumer protections cannot be enforced effectively when the role of AI is hidden or obscured. But narrow definitions and overbroad exemptions will render ADS regulations ineffective—as appears to have happened with New York City's law.

The risks associated with hidden ADSs are particularly high for disabled workers and consumers because the systems are often inaccessible or biased against people with disabilities. In the cases where consumers and workers have to interact with the system, as is the case with many automated employment assessments, the systems are often inaccessible and offer people with disabilities few or no options for accommodation or alternative assessment methods.

CDT and allied organizations have found extensive discriminatory impact caused by ADSs. A 2020 CDT report highlighted how algorithmic hiring tools can discriminate against disabled job candidates, noting that "as these algorithms have spread in adoption, so, too, has the risk of discrimination written invisibly into their codes." Last fall, CDT, the American Association of People with Disabilities (AAPD), and Coworker.org released a report, *Screened Out: The Impact of Digitized Hiring Assessments on Disabled Workers*, detailing the findings of a study on disabled workers' experiences with modern digitized assessments, including automated video game assessments and video interviews. The disabled workers who participated said that they "felt discriminated against and believed the assessments presented a variety of accessibility barriers." Notably, HB 1331 includes an unqualified exemption for "artificial intelligence-enabled video games," which would seem to exclude the sorts of gamified assessments that many study participants found inaccessible and discriminatory.

Other CDT publications have highlighted how electronic surveillance and algorithmic management (or "bossware") systems are used in ways that can violate the rights and threaten the health and safety of disabled workers, how surveillance technology can discriminate against disabled people, and how tenant screening algorithms can disproportionately exclude disabled

people, among other marginalized groups. The Disability Rights Education and Defense Fund (DREDF) published a brief in 2023 describing how algorithms could embed ableist standards of care into health care decision-making, expressing concern that "algorithmic and AI bias can further stigmatize patients, misdirect resources, and reinforce or ignore barriers to care rather than serving as a pathway to improving treatment and health outcomes."

These biases and barriers to accessibility cannot be addressed without basic transparency regarding when and how automated decision systems are being used, what those systems measure or assess, and how they measure or assess it. Systems that are unreliable, inaccurate, or biased against individuals with disabilities or people from other marginalized communities likewise will not be detected unless companies conduct impact assessments to identify potential sources of inaccessibility, bias, and invalidity.

What the bill covers (and what it does not)

HB 1331's notice, explanation, and accountability provisions focus on and are carefully tailored to apply only to ADSs that make or could alter the outcome of "consequential decisions" about consumers or workers—meaning decisions that have "legal or similarly significant effects" on a consumer's or worker's access to, cost of, or terms of major economic and life opportunities, specifically employment, education, financial or lending services, health care, and housing. In other words, **the law only applies to ADSs that are used to make certain key decisions that have major impacts on ordinary people's lives. The bill's requirements do not apply to any other AI systems, including generative AI systems like ChatGPT or any other forms of AI** unless the system is used in making one of those key decisions.

HB 1331 would arm consumers and workers with sorely needed information about life-altering decisions

Before using an ADS in a consequential decision, HB 1331 would require companies to provide the subject of the decision with information regarding the AI decision system, including "the purpose of the … system in use, including the nature of a decision … that is made…"

The bill does *not* require companies to disclose source code, training data, or trade secrets. In fact, many developers already routinely share the types of information HB 1331 would require on their websites and in their marketing materials and client pitches, and. That makes sense—after all, most deployers of AI decision systems want to understand how and why those systems work before purchasing them. HB 1331 would simply ensure that affected consumers and workers receive rudimentary information about those systems as well.

Providing this basic transparency would help resolve the dystopian situation that consumers and workers increasingly face, where their lives and careers are influenced and sometimes dramatically altered by algorithms they do not understand and, in many cases, do not even know exist. The Fair Credit Reporting Act has ensured consumers receive similar transparency for credit decisions for more than 50 years, starting soon after credit scoring started to become an automated process. The time has come for consumers to receive the same protection for today's increasingly automated decisions in other major spheres of life.

Why a broad definition of covered systems is essential

In other states considering similar legislation, legislators have been pressured to insert loopholes into two key definitions that define the bill's scope—"high-risk artificial intelligence system" and the related term "substantial factor." Specifically, some of the other bills' definitions of "high-risk AI system" are limited to systems that are "specifically intended" to make consequential decisions and/or have "human review" loopholes in the definition of "substantial factor." **Either of these loopholes would render the bill's protections largely meaningless; they should be avoided.**

An "intent" requirement would give companies a get-out-of-law-free card. A company could simply claim in technical or other documentation that their system is "not intended to overrule human decision making," and voilà—it could avoid the consumer protections built into the law and keep the system's existence hidden.

Similarly, exempting systems so long as there is "human review" would allow companies to evade the law even if the human "reviewer" is effectively a rubber stamp for the AI system in the decision process. These possibilities are not theoretical; already, companies have been repeatedly caught using AI decision systems without human review (as detailed in Hilke Schellmann's 2023 book, *The Algorithm*) or pressuring human "reviewers" to act as rubber stamps while telling the outside world that humans remained in control of the decision process. Moreover, research has shown that even well-intentioned humans tend to over-rely on algorithmic recommendations.

Because consumers and workers often have no idea when companies use ADSs on them, disclosure obligations can only be effective if a bill covers all systems that influence the decision-making process. Otherwise, a company could unilaterally evade the law—or human reviewers could, due to pressure or a desire to take the path of least resistance, act as rubber stamps—and the consumers and workers whose lives are altered as a result would be none the wiser. An effective law must avoid those pitfalls.

Changes needed to ensure the bill's effects match its intent

Replacing or narrowing the substantial factor requirement

The bill's definition of "high-risk artificial intelligence system" extends only to systems that make, or are a "substantial factor" in making, decisions with "legal effects concerning the consumer." This limit is reasonable on the surface but, for the reasons described in the preceding section, limitations based on the role that an AI system plays in a particular decision are prone to abuse because affected consumers, regulators, and other outsiders usually have no way of knowing how "substantially" a company is using an AI decision system (or, indeed, whether they are using the system at all).

Indeed, New York City's 2021 AI-in-hiring ordinance failed in large part because of a similar definition, which limited the bill to systems that "substantially assist or replace" human decision-making. The term "substantially assist" was undefined in the ordinance and interpreted

very narrowly by the city's enforcement agency. The narrow definition, combined with weak enforcement (addressed below), made it easy for deployers to effectively opt out of the law, as detailed in a study by researchers from Cornell University, Consumer Reports, and Data & Society.

We recommend eliminating the "substantial factor" requirement and defining "high-risk automated decision system" consistent with California Assembly Bill 2885, which was enacted last year and applies to any system "that is used to assist or replace" human decision-making with respect to a consequential decision. At a minimum, "substantial factor" should be explicitly defined in the law in a manner that clearly covers any system that is capable of altering the outcome of a decision or that serves as a basis or partial basis for a decision.

Adding a right to explanation

While HB 1331 includes pre-decision notice requirements and requires deployers to provide consumers with an opportunity to appeal and correct adverse decisions, it does not include a right to explanation. Without an explanation, consumers will not know whether there is anything to correct or appeal.

To ensure consumers can fully exercise their rights to correct and appeal, the bill should give consumers a right to specific and accurate explanations of why ADSs generated the output used in a consequential decision after the decision has been made. The explanation must be clear enough to ensure that consumers understand the output, what information it is based on, and whether it is based in whole or in part on inaccurate or improper information or inferences.

Such explanations are not a new idea. Congress adopted very similar rules for consumer finance decades ago during the earliest wave of automated decision-making. In the 1960s, banks began using computer-generated credit scores to decide which consumers could get loans and what interest rates each would pay. It became clear that many consumers were being denied credit without any explanation or based on inaccurate information, and consumers were outraged. Congress responded by passing the Fair Credit Reporting Act (FCRA) in 1970, giving consumers the right to know what factors go into credit decisions and a right to review and correct the underlying information that those decisions were based on.

It makes sense for companies to give consumers and workers receive similar transparency for AI-driven decisions so that they know when companies use AI to make key decisions about them, what factors those decisions are based on, and how they can correct inaccurate information.

Recommendation: Add the following language as a new section to the bill:

> *(A) On and after October 1, 2025, a deployer that has deployed a high-risk artificial intelligence system to make, or be a substantial factor in making, a consequential decision concerning a consumer shall provide to the consumer a single notice within 24 hours that includes:*

(1) A specific and accurate explanation that identifies the principal factors and variables that led to the consequential decision, the source(s) of data processed in the profiling decision, and a plain language explanation of how the consumer's personal data informed these principal factors and variables when the high-risk artificial intelligence system made, or was a substantial factor in making, the consequential decision;
(2) The role that the high-risk artificial intelligence system played in the decision-making process;
(3) Information on how the consumer can exercise their rights under 14–5004(A)(5) of this subtitle with respect to any personal data processed by the high-risk artificial intelligence system;
(4) Information about consumers' right to correct, and how the consumer can submit corrections and provide supplementary information relevant to the consequential decision; and
(5) What actions, if any, the consumer might have taken to secure a different decision and the actions that the consumer might take to secure a different decision in the future;

(B) No deployer shall use a high-risk artificial intelligence system to make, or assist in making, a consequential decision if it cannot provide a specific and accurate notice that satisfies the requirements of this section and section 14–5004.

Clarifying pre-decision notice requirements

The bill requires deployers to provide impacted consumers with a pre-evaluation notice that includes "a plain language description of" the system. This is a good start, but it would not provide consumers or workers with crucial information about what AI decision systems are supposed to measure and how they measure it. At a minimum, consumers and workers should know what types of decisions AI systems are used to make, what personal data and attributes are considered, and how systems use that information to generate outputs. That information is essential for consumers and workers to make an informed choice about whether to subject themselves to such an assessment. In the employment context, robust notices are particularly critical for individuals with disabilities, who otherwise will not know whether they may need to seek accommodation.

Recommendation: Expand pre-evaluation notice requirements to give individuals the right to know what information the system uses and how the system uses that information to make a decision:

- *[Deployer must provide the consumer with] a description, in plain language, of such high-risk artificial intelligence system, which description shall, at a minimum, include:*
  (i) The personal characteristics or attributes that the system will measure or assess;

(ii) The method by which the system measures or assesses those attributes or characteristics;
(iii) How those attributes or characteristics are relevant to the consequential decisions for which the system should be used;
(iv) Any human components of such system;
(v) How any automated components of such system are used to inform such consequential decision; and
(vi) A direct link to a publicly accessible page on the deployer's website that contains a plain-language description of the logic used in the system, including the key parameters that affect the output of the system; the system's outputs; the type(s) and source(s) of personal information collected from natural persons and processed by the system when it is used to make, or assists in making, a consequential decision; and the results of the most recent impact assessment, or an active link to a webpage where a consumer can review those results.

- Define "personal information" as: "any information, including derived data and unique identifiers, that is linked or reasonably linkable, alone or in combination with other information, to an identified or identifiable individual or a device that identifies or is linked or reasonably linkable to an individual."

Eliminate or significantly narrow the exemption for trade secrets and "security risks"

The bill completely exempts information developers consider a "trade secret" or "a security risk." These exemptions essentially require total deference to developers in deciding for themselves what information is a "trade secret" or a "security risk." Companies frequently over-designate ordinary business information as a trade secret. Infamously, Theranos asserted that information showing its blood-testing technology did not work was a "trade secret"—a fact that only came to light after the company's fraudulent scheme fell apart. We would not want a similar situation with information about ADSs that cause discrimination being kept secret. Moreover, unscrupulous companies could seize on the undefined term "security risk" to justify, for example, intrusive use of ADSs to monitor, discipline, and terminate workers.

Indeed, these exemptions are unnecessary given the modest nature of developers' disclosure requirements, all of which is information that responsible developers already provide to potential customers quite readily. The bill does not call for any company to disclose its source code, training data, or any other secret sauce that would undermine legitimate intellectual property rights. It simply calls for developers to provide deployers with the information necessary to understand, operate, and manage the system and understand its limitations. A trade secret or confidentiality claim regarding such information would be doubtful at best. That interest should, in any event, yield to the interest in ensuring that deployers have the information they need to comply with their obligations under the law and monitor the performance of systems used in consequential decisions.

Recommendation: These exemptions should be eliminated. If that cannot be done, they should be narrowed—companies should be permitted to redact information only pursuant to trade

secret law, and when they do, they should be required to alert deployers or other audiences where and why they redacted such information.

Strengthening enforcement

Many laws purporting to protect workers, consumers, or the public at large have failed to accomplish their goals because they lacked strong enforcement provisions. If a company can write off potential fines, penalties, or other consequences under a law as a mere inconvenience or as the cost of doing business, then companies may simply ignore that law.

The current enforcement provisions threaten a similar fate for HB 1331. The current text would make violations of the act an "unfair, abusive, or deceptive trade practice" under Maryland's Consumer Protection Act, but it then strips away the most effective enforcement mechanism: a consumer's right to pursue a civil action. This limitation will likely result in an overburdened AG's office and an underenforced law.

Consumers harmed by violations of HB 1331 should have access to all remedies that Maryland's consumer protection laws allow, including a private right of action. Individual consumers are in the best position to know when a company has used an AI tool improperly or unfairly to make an important decision about their lives, and they should have the ability to take action to vindicate their rights under this law.

Recommendation: Delete the words "except for § 13–408 of this article" on page 13, line 22.

Conclusion

HB 1331 represents a critical step forward in protecting workers and consumers from the risks associated with often-hidden automated decision-making technologies. By addressing key issues such as notice and accountability, the bill would empower workers and consumers and ensure that they have a basic understanding of systems that increasingly alter the course of our lives. We suggest modest revisions to ensure consumers receive adequate information about AI systems used in life-altering decisions and to ensure the bill's provisions provide the intended protection for Marylanders.

HB 1331 has the potential to set a strong precedent for safeguarding consumers' and workers' rights in Maryland and beyond. Thank you again to Delegate Qi for your excellent work on this bill. We look forward to continuing to engage as it moves through the legislative process.