**Testimony of Eric Null, Co-Director, Privacy & Data Program,**
**Center for Democracy & Technology**
**Before the Maryland Senate Finance Committee**
**February 27, 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 greatly appreciate the effort that Senator Hester and her colleagues have put into crafting SB 936, which would cover automated decision systems (ADSs)[1] in settings as diverse as employment, housing, health care, and criminal justice. We are generally supportive of the transparency provisions in the bill, as they require transparency around key aspects of ADSs. We are also generally supportive of the duty of care placed on developers to take steps to affirmatively prevent algorithmic discrimination.

Unfortunately, as it stands, the bill's loopholes and exemptions would undermine the bill's goals because they essentially allow companies to opt themselves out of complying with the law. Our concern is not hypothetical: companies have almost entirely ignored New York City's AI-hiring ordinance, which has similar infirmities.

Below, we briefly discuss the problem of companies using hidden algorithms and other ADSs to make key decisions about workers and consumers and how the lack of transparency surrounding such systems poses particular risks to disabled workers and consumers. We then highlight the bill's strong transparency provisions, but discuss how, in its current form, the bill's narrow definitions and overbroad exemptions would make those transparency provisions largely illusory.

We hope that the Committee will take steps to address these issues by fixing the key definitions and closing these loopholes while there is still time to do so. If those changes are made, we would support SB 936.

---

[1] We use "automated decision system" to refer to the class of technologies that are the subject of SB 936. We do not necessarily oppose the use of the term "high-risk artificial intelligence system" for purposes of this bill, but that term may be confusing since it has also been used to refer to a wide range of other AI applications that also pose risks to large numbers of consumers (such as AI-enabled surveillance systems, deepfakes, etc).

**Automated decision systems are frequently hidden from the workers and consumers they effect 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 and AAPD 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, SB 936 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 discriminates 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.

## **We support the substantive transparency requirements in SB 936**

The bill's substantive provisions would grant consumers the right to basic transparency regarding ADSs. The bill's notice provisions would require covered companies to provide consumers with:

> (a) a description of the personal characteristics or attributes that such system will measure or assess, (b) the method by which the system measures or assesses such attributes or characteristics, (c) how such attributes or characteristics are relevant to the consequential decisions for which the system should be used, (d) any human components of such system, and (e) how any automated components of such system are used to inform such consequential decisions.

Consumers and workers deserve this information as a matter of basic fairness when they are subjected to important decisions that could affect, for instance, their housing, career compensation, health, safety, or economic security. These disclosures are especially crucial for people with disabilities, many of whom need the information to know whether the system might be inaccessible or inaccurate with respect to them and whether the company's human-review

safeguards provide an adequate opportunity for them to request accommodation or an alternative form of evaluation or decision process.

## SB 936's definitions would leave workers and consumers unprotected unless they are amended

### *The definition of "consumer" excludes workers and could be misused to exclude many consumers*

ADSs are being aggressively marketed and implemented in workplaces and labor markets in Maryland and across the country. But the bill's definition of "consumer" explicitly excludes anyone acting "in a commercial or employment context." This broad and undefined exemption would apparently exclude all workers—employees, independent contractors, and job candidates alike—from the bill's protections.[2] This means that, despite the fact that the bill contemplates "access to or provision of employment" as a "consequential decision," those protections would not apply to any Maryland worker. In other words, the bill would provide no protection when employers use inaccessible AI systems or flawed or biased algorithms to decide who to hire and fire and how much workers get paid. Additionally, the term "commercial . . . context" is ambiguous and undefined; for instance, are customers browsing in a retail store acting in a "commercial . . . context" and thus any use of ADSs in that store, such as facial recognition technology that seeks to identify shoplifters, not subject to the bill? The definition of "consumer" thus should be amended to cover all Marylanders.

**Recommendation**: The definition of "consumer" should simply be: "a natural person who is a resident of this state."

### *The bill's definition of "high-risk artificial intelligence system" would make it far too easy for companies to evade the law's transparency requirements*

SB 936 defines "high-risk artificial intelligence systems" as systems that are "specifically intended to autonomously make, or be a substantial factor in making" key decisions. This definition is deeply concerning; it would make it far too easy for developers to define themselves out of the law's scope. A company could, for example, completely avoid the bill's obligations—including its obligation to reveal the existence of the system to affected consumers—by crafting technical documentation or marketing materials with a disclaimer saying: "Use only with human supervision" or "Not intended to serve as the principal basis for a decision." The developer and any deployer that uses the system could then justify to themselves that the ADS is not intended to autonomously make or substantially affect decisions and use that as a justification to ignore the law. It would be almost impossible for a plaintiff to disprove the developer's subjective intent—especially because invoking this definitional loophole would allow companies to keep the system's very existence hidden.

This loophole thus would allow companies to opt themselves out of complying with the law by determining its system is exempt from the law's disclosure obligations. As a result,

[2]

consumers, outside experts, and enforcers would have no way to learn about the system and challenge that conclusion.

**Recommendation**: The definition should be amended to eliminate any requirement tied to the developer's or deployer's subjective intent, and should focus on the system's use.

### *The definitions of "substantial factor" and "principal basis" could allow companies to escape transparency and accountability by having humans rubber-stamp algorithmic "recommendations"*

The definitions of "substantial factor" and "principal basis," as currently written, would allow companies using ADSs to evade the law if there is "human review, oversight, involvement, or intervention" in the decision process or "meaningful consideration by a human" in the decision process. Under this definition, as long as a company assigns a human to review algorithmic "recommendations," the company could justify to itself that it is exempt, keep the system hidden from affected consumers, and ignore the law's transparency and accountability requirements.

The problem is that companies often claim they have human review for all decisions—but, in reality, those humans can act as rubber stamps or feel pressure to follow algorithmic recommendations, exemplified by the investigative reports on Cigna described above. Thus, as with the "specifically intended" loophole, this provision effectively would provide companies (in this case deployers) an easy avenue to opt out of compliance with the law.

In a brief on disability bias in health care algorithms, DREDF highlighted how the incorporation of algorithms into decision-making processes can lead humans to abdicate their responsibilities even when there is a genuine desire for humans to bring their own judgment to bear as well:

> While DREDF is concerned with how algorithms are created and how developers evaluate the fairness of the formula and data inputs used, the crux of our concern with computer-mediated tools is that the human decision-makers who bear ethical and professional obligations as health care providers and entities have changed their decision-making process. Furthermore, they may choose to do so without any notice of the change. In essence, they may believe they have fully delegated their decision-making authority and should no longer be held accountable for the discriminatory outcomes because computers cannot "intend" discrimination. Once algorithms are involved and assigned a role within decision-making, there is a human tendency to give primary weight to the algorithmic output, decision, or recommendation, even in the face of conflict with human expertise, knowledge, and judgment. Examples of this deference to algorithms can be found in decisions made by pilots who defer to automatic flight control systems, as well as by physicians making treatment decisions in critical care units; the higher the stakes and, some might say, the greater the need for a human grappling with ethics, life values, and implicit bias, the greater the pressure to abdicate responsibility to an "objective" algorithm.

Having a "human in the loop" thus is no panacea. "Substantial factor" requirements and, especially, "human review" exemptions instead introduce loopholes with the potential to dramatically undermine the law.

**Recommendation**: The "substantial factor" requirement should be eliminated and the definition of "high-risk artificial intelligence system" amended to align with the definition of "high-risk automated decision system" in California Assembly Bill 2885, which was enacted last year and applies to all systems that "assist or replace human discretionary decisions." That definition would ensure that affected consumers receive notice of any automated decision system used in crucial decisions about them.

### *SB 936 contains other exemptions and carve-outs that would allow too many companies to avoid compliance, transparency, and accountability*

The bill contains several other exemptions that would significantly undermine the bill's substantive provisions unless they are eliminated.

SB 936 gives companies discretion to ignore disclosure requirements by designating information as "confidential" or "proprietary"

The bill completely exempts information developers consider a "trade secret," "confidential or proprietary information," or "a security risk." These exemptions essentially require total deference to developers in deciding for themselves what information is a "trade secret," "confidential," "proprietary," or a "security risk."[3] Companies frequently over-designate ordinary business information as a trade secret, or confidential or proprietary when it suits them to do so. 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.

Indeed, the exemption is 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

[3] While "trade secret" is defined in Maryland's laws here, that definition still gives companies leeway to determine what information has value by being kept secret.

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

The definition of "high-risk artificial intelligence system" should be clarified so that it covers all uses of AI in consequential decisions

The definition of "high-risk artificial intelligence system" contains several additional exemptions that carve entire industries and products out of the bill's scope. For example, the bill exempts "autonomous vehicles" (AVs). Thus, developers and deployers of AV-related ADSs would be under no obligation to, for instance, protect against known or reasonably foreseeable risks of algorithmic discrimination.

Yet, while future advances in AVs could improve accessibility and lower the cost of travel, there are instances in which decision systems embedded in AVs could pose a significant risk of discrimination. AVs have been known to make decisions on the road that ignored wheelchair users, people who lack limbs, or people with other disabilities that cause them to look or act "atypically." If AVs are excluded, there would be no protection if an ADS embedded within an AV was trained or learned to pass by humans who have service animals (a common behavior among human drivers) or give priority to traffic and fail to stop at the curb cuts needed by a wheelchair user.

Similarly, the definition of "high-risk artificial intelligence system" exempts "artificial intelligence-enabled video games." Although the intent may be for this exemption to apply to games used for entertainment, it is problematic as written because AI-enabled video games are a common form of ADS in hiring, and such "games" are often inaccessible to people with disabilities or designed in ways that introduce bias.

**Recommendation**: All exemptions should either be eliminated or should apply only where ADSs are not used in making a consequential decision.

Exemptions for specific industries should be narrowed

Other exemptions appear to carve out some of the industries that the bill's definition of "consequential decisions" purports to cover. For example, even though the definition of "consequential decisions" includes financial services, lending, and insurance, § 14-47A-02(2) contains a carve-out exempting insurers subject to the Maryland Insurance Administration (MIA). Insurance companies should have to comply with the bill's provisions—in particular, where they do not overlap with, for instance, MIA protections for AI use. Where the requirements are substantially similar, compliance with one should be sufficient for both.

Likewise, § 14-47A-02(1) exempts any system "acquired by" the federal government, and excludes *any* system (except ADSs making employment or housing decisions) that the federal government *ever* acquires—even when those systems are being used by private companies. The implicit justification for this section, that there are additional protections at the federal level for ADSs procured by the federal government, is not always accurate. Further, federal government use of ADSs is already exempt because "person" does not include "government units."

**Recommendation**: Exemptions for certain types of ADSs subject to other authorities should apply only where the requirements of this bill and the other authorities are substantially similar. The exemption for ADSs acquired by the federal government should be removed.

Exemption for companies complying with "risk management frameworks" should be removed

The bill also contains an exemption allowing companies to be "presumed to be in conformity" with the bill if they are "in conformity with the latest version of the Artificial Intelligence Risk Management Framework published by the National Institute of Standards and Technology, Standard ISO/IEC 42001 of the International Organization for Standardization, or another nationally or internationally recognized risk management framework." But none of these standards presently comes close to ensuring that consumers receive the information that the bill's substantive transparency provisions would provide, and it is impossible to predict whether they will be strengthened, weakened, or eliminated as administrations and the composition of standards bodies change. Moreover, the words "another nationally or internationally recognized risk management framework" appear to give companies wide discretion to pick and choose the standards of their liking. This will give rise to inconsistencies and would also make enforcement difficult.

**Recommendation**: Eliminate the presumption of compliance for companies that conform with non-statutory risk management frameworks.

Companies could escape the right to appeal requirement by invoking a vague "best interest of the consumer" exemption

Individual requirements also contain some loopholes that would allow companies to escape the bill's obligations. For example, a developer could take away an individual's right to appeal a decision made by an automated system—an important protection that this bill would rightly give consumers—if the company decides that it is not in "the best interest of the consumer" to let that person appeal the decision. The law does not define what constitutes "the best interest of the consumer;" conceivably, for example, an auto insurance company could say that an appeal of an ADS-driven decision setting a high premium for a consumer is not in the consumer's best interest because it would delay the consumer's ability to obtain coverage. Further, the law puzzlingly allows the company in control of the AI system to decide what is in the consumer's best interest rather than the consumer, who could simply choose not to appeal the decision.

**Recommendation**: The "best interests of the consumer" exemption should be removed. At the very least, the "best interests" exemption should be significantly clarified and narrowed, much like the right of appeal's other exemption, which allows appeal denial if there is a delay that poses a risk to the life or safety of a consumer.

## Accessibility and accommodation decisions covered by Maryland's human rights laws should be covered consequential decisions

We recommend that the bill's list of consequential decisions be amended to incorporate decisions that grant or deny disabled Marylanders' accessibility and accommodation rights. Maryland has notably provided its residents with some of the country's strongest disability rights protections, with stronger accessibility and accommodation requirements in many cases than federal law provides. But the current list of consequential decisions does not appear to cover, for example, a business's use of ADSs that provide language interpretation and translation services in health care, even though a poorly designed system could easily have the practical effect of denying disabled Marylanders' access to health care.

**Recommendation**: Add the following to the list of consequential decisions: "(x) a reasonable accommodation or other right granted under the civil rights laws of this state."

## Conclusion

While SB 936 aims to address critical risks associated with ADSs, improvements are needed to ensure the bill's provisions provide the intended protection for Marylanders. As it stands, the bill's narrow definitions and ambiguous and overbroad exemptions undermine its otherwise-strong transparency and accountability provisions. To truly safeguard the rights of Marylanders, including and especially those from vulnerable groups, the bill should be amended to close these loopholes and strengthen key definitions. If those changes are made, Maryland will have an opportunity to lead the country with strong legislation empowering its workers and consumers and promoting trustworthy and rights-protecting innovation.