

Testimony of Julia Stoyanovich on Regulation of High-Risk Use of Artificial Intelligence in the Private Sector

Hearing of the New York State Senate Standing Committee on Internet and Technology

January 15, 2026

Note: ChatGPT 5.2 was used for stylistic purposes when drafting this testimony.

Thank you, Chair Gonzalez, and members of the Senate Standing Committee on Internet and Technology, for the opportunity to submit my testimony on the regulation of high-risk uses of artificial intelligence (AI) in the private sector.

AI regulation is more important now than at any prior moment. Systems with extraordinary capabilities are being deployed at speed across the economy, often without meaningful oversight. In recent weeks, Grok, a widely used AI system, publicly generated non-consensual sexual images and extremist content—not because of a one-off bug, but because boundary-pushing drove engagement and profit. That episode was not an anomaly; it was a warning. Across sectors, we are already seeing AI systems that fabricate authoritative-sounding information¹, deny people jobs or healthcare coverage based on opaque and unstable criteria², and quietly shape access to housing, credit, and education³. At the same time, meaningful guardrails at the federal level remain limited and uncertain.

This is precisely the risk profile that The New York AI Act (S.1169-A) is designed to address: AI systems that are a substantial factor in consequential decisions or that materially impact individuals' rights, liberties, safety, or welfare. This combination creates not only risk, but urgency. How New York responds matters.

My main point today is simple: Finding a clear, enforceable, and forward-looking approach to regulating high-risk AI is not a drag on innovation; it is an economic development advantage.

Regulatory clarity attracts responsible investment, protects businesses from arbitrary and unsafe systems, and positions New York as a leader at what has rightly been described as a modern-day Sputnik moment for AI.⁴ Waiting is not neutral. Acting decisively now allows New York to shape the incentives, markets, and norms around AI in a way that

¹ <https://apnews.com/article/new-york-city-chatbot-misinformation-6ebc71db5b770b9969c906a7ee4fae21>

² <https://www.washingtonpost.com/business/2025/12/01/ai-work-regulations-california>

³ <https://www.theverge.com/news/801205/new-york-rent-price-fixing-ban-software>

⁴ Stoyanovich, "Deep Seek: A Deep Dive," Testimony at the Hearing of the Committee on Science, Space and Technology of the U.S. House of Representatives, Research and Technology Subcommittee, April 2025, <https://science.house.gov/2025/4/deepseek-a-deep-dive>

both protects New Yorkers and strengthens the state's long-term competitiveness.

By way of introduction, I am an Associate Professor of Computer Science & Engineering and of Data Science, and the founding Director of the Center for Responsible AI at New York University. My research focuses on AI and data engineering systems, with an emphasis on incorporating legal requirements and ethical norms into how these systems are designed, developed, and used.⁵ I teach responsible AI to students, practitioners in industry and government, and members of the public⁶. I have been deeply involved in AI governance and regulation in New York City, New York State, nationally, and internationally since 2017⁷. I am a recipient of the Presidential Early Career Award for Scientists and Engineers (PECASE), by nomination from the National Science Foundation. I have provided testimony across multiple legislative and regulatory venues, and I will draw on that experience in my remarks today.

I strongly support the goal of the proposed New York AI Act: to protect New Yorkers from the adverse impacts of high-risk AI systems. This regulation is both timely and necessary. S.1169-A is best understood as complementary to the RAISE Act, addressing the everyday, high-risk uses of AI systems that already shape employment, housing, credit, healthcare, and access to essential services. By focusing on deployer obligations, downstream use, civil rights protections, and market integrity, the bill fills a critical gap left by frontier-model legislation. Together, these approaches support a layered governance framework that addresses both extreme risks at the technological frontier and routine but consequential harms at scale.

I organize my remarks around five core points that are essential for effective, durable, and pro-innovation AI governance.

- **[Recommendation 1:](#)** Support the bill's focus on high-risk uses of AI, and explain why this scope protects New Yorkers and provides regulatory clarity to businesses.
- **[Recommendation 2:](#)** Preserve and strengthen the notice and opt-out provisions, arguing that transparency and individual agency are key enablers of trust and adoption.
- **[Recommendation 3:](#)** Emphasize the importance of assessing validity and reliability as economic, not merely ethical, concerns.
- **[Recommendation 4:](#)** Treat data protection and data governance as integral to AI regulation and long-term innovation.
- **[Recommendation 5:](#)** Highlight the need for education, expertise, and institutional capacity within government to ensure consistent enforcement that supports responsible economic development.

⁵<https://r-ai.co/>

⁶ <http://r-ai.co/education>

⁷ <http://r-ai.co/policy>

Recommendation 1: Maintain the Bill's Focus on High-Risk Uses of AI

S.1169-A grounds its scope in clear statutory definitions of a high-risk AI system, defined as one that is a substantial factor in a consequential decision or that materially impacts individuals' rights, liberties, safety, or welfare (§ 85(10)), and of a consequential decision itself (§ 85(4)). These definitions are reinforced by the Legislature's findings, which recognize both the pervasiveness of AI in high-impact domains and the need for targeted safeguards rather than blanket regulation (§ 2). Together, these provisions establish a risk-based framework that is precise enough to be enforceable and flexible enough to endure technological change.

The bill's focus on AI systems that materially affect rights, opportunities, safety, and welfare, rather than attempting to regulate AI wholesale, is both legally defensible and practically workable. Legally, it aligns regulatory obligations with foreseeable harm, a well-established principle in civil rights and consumer protection law. Practically, it avoids overregulating low-risk or experimental uses while ensuring that systems deployed in hiring, housing, credit, healthcare, education, and similar contexts are subject to meaningful oversight. These are the settings where AI errors and abuses have the greatest human and economic consequences, and where regulation is most clearly justified.

Why this strengthens competitiveness. Risk-based regulation provides clarity to developers and deployers about when heightened obligations apply. That clarity is a competitive asset. Companies can innovate freely in low-risk settings, including prototyping, internal tools, and benign consumer applications, while investing responsibly in safeguards for high-stakes uses. This predictability lowers compliance uncertainty, reduces the risk of retroactive enforcement, and encourages capital investment in jurisdictions with clear rules. New York, in particular, has benefited when regulation has reduced ambiguity rather than eliminated risk.⁸

History offers strong, New York-relevant parallels. In financial services, risk-differentiated regulation helped make New York a global fintech hub by clarifying which activities required heightened oversight and which could scale more freely.⁹ Similarly, risk-based data protection regimes, including the GDPR, catalyzed markets for privacy-enhancing technologies and compliance services, many of which now have major operations in New York.¹⁰ More recently, New York State agencies have aligned with the NIST AI Risk Management Framework, which

⁸ NYDFS Cybersecurity Regulation, 23 NYCRR 500:

https://www.dfs.ny.gov/industry_guidance/cybersecurity

NYDFS Cybersecurity Regulation: *Frequently Asked Questions and Industry Guidance.*

https://www.dfs.ny.gov/industry_guidance/cybersecurity/faqs

⁹ Basel III framework: <https://www.bis.org/bcbs/basel3.htm>

¹⁰ GDPR, Article 24: <https://gdpr.eu/article-24-responsibility-of-the-controller/>

explicitly adopts a risk-tiered approach to reduce uncertainty for innovators while focusing on safeguards where harms are most likely to occur.¹¹ In each case, proactive differentiation by risk did not slow innovation. It directed innovation toward durable, trustworthy products and reinforced New York’s position as a place where regulated industries can grow with confidence.

Background from my prior work. My work on responsible data engineering and AI emphasizes that most real harms do not arise from experimental systems at the technological frontier, but from routine deployment in consequential contexts.¹² In hiring, for example, tools that appear mundane, such as resume screeners or personality assessments, can quietly produce arbitrary or discriminatory outcomes at scale.¹³ Overbroad or vague regulation exacerbates this problem by creating uncertainty for responsible actors while allowing harmful practices to persist in gray areas. By contrast, narrowly scoped regulation creates stable markets: it tells companies where scrutiny is expected, incentivizes investment in testing and documentation, and surfaces failures before they scale.¹⁴ This is precisely the function a risk-based framework is meant to serve.

Comparison to the RAISE Act. The RAISE Act focuses on frontier models and catastrophic-risk scenarios—an important but narrow slice of the AI risk landscape. S.1169-A fills a critical gap by addressing the everyday, high-risk uses of AI that already shape New Yorkers’ lives. These are not speculative future harms; they are present-day market realities. By concentrating on deployer obligations and downstream use in consequential settings, S.1169-A complements frontier-focused legislation and completes the governance picture. Together, the two approaches form a layered framework: one that mitigates extreme risks at the frontier while ensuring accountability and market integrity where AI is most commonly deployed today.

Recommendation 2: Preserve and Strengthen Notice and Opt-Out Requirements

S.1169-A includes robust notice, opt-out, and appeal provisions for high-risk AI systems used in consequential decisions (§ 86-a(1)–(2)). These provisions require deployers to inform individuals in advance when AI systems will be used, to provide a meaningful opportunity to opt out of automated decision-making in favor of human review, and to notify individuals after a decision is made of their right to contest and appeal that decision. Taken together, these

¹¹ NIST AI RMF 1.0: <https://www.nist.gov/itl/ai-risk-management-framework>

¹² Stoyanovich and Howe., “Follow the Data! Algorithmic Transparency Starts with Data Transparency,” New America, <https://www.newamerica.org/pit/blog/follow-data-algorithmic-transparency-starts-data-transparency/>

¹³ Rhea et al., “Resume Format, LinkedIn URLs and Other Unexpected Influences on AI Personality Prediction in Hiring,” AIES 2022, <https://dl.acm.org/doi/10.1145/3514094.3534189>

¹⁴ NIST AI Risk Management Framework 1.0, <https://www.nist.gov/itl/ai-risk-management-framework>

requirements establish a baseline of transparency and agency that is essential for lawful, accountable AI deployment in high-stakes contexts.

These provisions are among the most important in the bill. Notice and opt-out directly address the severe information asymmetries that characterize AI-mediated decision-making. Individuals subject to automated decisions often do not know that AI is involved, what role it played, or what criteria were used. Without notice, people cannot meaningfully exercise their rights, seek accommodations, or contest errors. Without opt-out and appeal, automated systems risk becoming unreviewable gatekeepers. By contrast, S.1169-A ensures that AI remains a decision-support tool rather than an unaccountable authority.

Why this strengthens competitiveness. Notice and opt-out requirements are sometimes seen as burdensome, but in practice they are low-cost for deployers and high-value for markets. Transparency builds trust, and trust is a prerequisite for adoption, particularly in high-stakes, AI-mediated decision-making.¹⁵ When individuals understand how AI is used and retain agency, they are more willing to engage with AI-enabled services, which benefits responsible companies by reducing reputational risk, increasing user confidence, and stabilizing demand.¹⁶ Conversely, markets characterized by opacity and surprise are brittle; they invite backlash, litigation, and abrupt regulatory intervention.¹⁷ By normalizing disclosure and human fallback, S.1169-A helps create a predictable environment in which AI products can scale.

There are clear parallels in other regulated domains that are directly relevant to New York. In consumer finance, disclosure and interoperability requirements, most notably the UK's Open Banking reforms, did not prevent innovation; they enabled it by allowing users to understand risk, compare products, and securely share data, which in turn catalyzed a wave of fintech innovation that has since shaped global markets and strongly influenced New York's own financial technology ecosystem.¹⁸ Similarly, in data protection, risk-based consent and transparency obligations under the GDPR catalyzed new markets in consent management, privacy engineering, and compliance-as-a-service, many of which now have major operations in New York.¹⁹ In each case, notice functioned not as friction, but as infrastructure. The same is true here. By establishing clear expectations around disclosure and choice, New York positions itself as a jurisdiction where AI-enabled services can grow sustainably and responsibly.

¹⁵ OECD, *Building Trust in AI*, <https://www.oecd.org/going-digital/ai/building-trust-in-artificial-intelligence>

¹⁶ Edelman Trust Institute, *Trust and Technology*, <https://www.edelman.com/trust/2024/trust-barometer>

¹⁷ FTC, *Aiming for Truth, Fairness, and Equity in Your Company's Use of AI*, <https://www.ftc.gov/business-guidance/blog/2021/04/aiming-truth-fairness-equity-your-companys-use-ai>

¹⁸ UK Open Banking Implementation Entity, <https://www.openbanking.org.uk/what-is-open-banking/impact-of-open-banking/>

¹⁹ European Commission, https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/data-protection_en

Background from my prior work. My research on algorithmic hiring illustrates why notice and opt-out are critical. Many hiring systems screen candidates based on proxies or inferred traits that job seekers are unaware of and cannot contest. I have proposed the use of simple, standardized “nutritional labels” for AI systems, short disclosures that explain what a system does, what data it uses, and what it is intended to optimize.^{20 21} See Figure 1 below for an example.²² These labels support informed consent, enable requests for accommodation, and surface validity concerns early. In practice, disclosure often reveals that systems rely on weak or unstable signals, prompting both candidates and employers to question their use. Notice is therefore not merely a procedural safeguard; it is a mechanism for improving system quality.

ACCOUNTANT	
Acme Partners	
Qualifications:	BS in accounting, GPA >3.0, Knowledge of financial and accounting systems and applications
Personal data to be analyzed:	An AI program could be used to review and analyze the applicant's personal data online, including LinkedIn profile, social media accounts and credit score.
Additional assessment:	AI-assisted personality scoring
ALERT: Applicants for this position DO NOT have the option to selectively decline use of AI analysis for any of their personal data or to review and challenge the results of such analysis.	

Figure 1: A posting label is a short, simple, and clear summary of the screening process. This label is presented to a job seeker before they apply, supporting informed consent, allowing them to opt out of components of the process or to request accommodations.

Comparison to the RAISE Act. The RAISE Act emphasizes transparency to regulators and reporting of safety incidents for frontier models, but it does not center the experience of individuals subject to AI-assisted decisions. S.1169-A fills this gap by establishing rights to notice, opt-out, and appeal at the point of impact. This distinction matters. Frontier-model transparency without end-user transparency leaves most everyday AI harms unaddressed. By contrast, S.1169-A ensures that transparency and accountability reach the people most directly affected. Together, the two approaches are complementary: RAISE focuses on oversight of extreme technical risk, while S.1169-A embeds agency and contestability into routine, high-risk uses of AI across the economy.

²⁰ Stoyanovich and Howe, “Nutritional Labels for Data and Models,” *Data Engineering Bulletin* (2019), <http://sites.computer.org/debull/A19sept/p13.pdf>

²¹ Stoyanovich *et al.*, “The Imperative of Interpretable Machines,” *Nature Machine Intelligence* (2020), <https://www.nature.com/articles/s42256-020-0171-8>

²² Stoyanovich, “Hiring and AI: Let job candidates know why they were rejected,” *The Wall Street Journal Reports: Leadership* (2021), <https://www.wsj.com/articles/hiring-job-candidates-ai-11632244313>

Recommendation 3: Explicitly Treat Validity and Reliability as Economic, Not Merely Ethical, Concerns

S.1169-A appropriately requires audits to assess not only risks of algorithmic discrimination, but also the accuracy and reliability of high-risk AI systems with respect to their intended and actual use cases (§ 87(2)(b)(ii)). These requirements are reinforced by the obligation for developers and deployers to establish and maintain risk management policies and programs that identify, document, and mitigate reasonably foreseeable risks over the lifecycle of a system (§ 89). Together, these provisions recognize that whether a system works as intended is a core component of responsible deployment.

The bill's expansion of auditing beyond bias and discrimination to include validity and reliability is an important and necessary step, but as currently framed it is not sufficient. Requiring audits to consider accuracy and reliability without clearer guidance on how validity should be assessed risks repeating a familiar pattern, in which formal compliance masks substantive failure. As I have argued in the context of automated hiring systems (NYC Local Law 144)²³, as well as in my Congressional testimony on high-impact AI²⁴, systems can satisfy narrow technical metrics while still failing in practice because they do not measure what they claim to measure, behave inconsistently under trivial input changes, or are poorly matched to their intended use. In consequential domains, even small instabilities or mismatches can scale into widespread harm. To avoid legitimizing systems that are unbiased but arbitrary, implementation of S.1169-A should more precisely specify expectations for assessing validity, including whether systems perform consistently, whether outputs are stable under task-irrelevant variations, and whether claims made by vendors about system capabilities are empirically justified. Without such specificity, audits risk becoming procedural rather than probative, undermining both the bill's protective goals and its promise of market discipline.

Why this strengthens competitiveness. Validity and reliability are not abstract ethical ideals; they are economic necessities. Invalid systems waste deployer resources, expose firms to legal and reputational risk, and distort markets by crowding out better tools with products that succeed through opacity rather than performance. When companies invest in AI systems that are arbitrary or unstable, they incur direct costs, including poor decision quality and missed opportunities, as well as indirect costs from downstream liability and loss of trust. By requiring evidence that high-risk systems work as intended, S.1169-A protects businesses from paying for tools that are ineffective or misleading, and helps level the playing field for vendors that invest in rigorous testing, documentation, and quality assurance. This kind of regulation channels innovation toward systems that are robust, dependable, and economically valuable.

²³ https://rules.cityofnewyork.us/wp-content/uploads/2022/12/Stoyanovich_144_Jan23_2023.pdf

²⁴ <https://www.schumer.senate.gov/imo/media/doc/Julia%20Stoyanovich%20-%20Statement.pdf>

Background from my prior work. My own research has documented dramatic validity failures in widely used AI tools, particularly in hiring. In audits of algorithmic personality assessments, my collaborators and I found that trivial, job-irrelevant changes to input format, such as submitting the same resume as a PDF versus plain text, could produce substantially different outputs for the same individual.²⁵ Systems with this level of instability cannot be considered valid measurement instruments. These failures harm workers by subjecting them to arbitrary decision-making, and they harm employers by undermining hiring outcomes and exposing firms to legal risk. Validity failures are therefore not a tradeoff between fairness and efficiency; they undermine both.

Comparison to the RAISE Act. The RAISE Act focuses primarily on the safety and security risks associated with frontier models, including catastrophic misuse and large-scale harm. While these concerns are important, they do not address a different and equally pervasive failure mode: AI systems that are confidently wrong. S.1169-A fills this gap by focusing on the everyday deployment of AI in consequential decisions, where systems may appear polished and authoritative while producing unreliable or unjustified outputs. By addressing validity and reliability directly, S.1169-A complements frontier-focused legislation and helps ensure that AI systems used throughout the economy are not only secure, but actually fit for purpose.

Recommendation 4: Treat Data Protection and Data Quality as Preconditions for Responsible AI Markets

S.1169-A appropriately integrates data considerations into its core risk management and reporting framework. The bill requires developers and deployers to account for the sensitivity and volume of data processed by high-risk AI systems as part of their risk management policies and programs (§ 89(1)(d)). It also mandates regular reporting and documentation that describe system design, training data, intended and actual uses, and known risks (§ 88). Together, these provisions recognize that AI risk cannot be assessed or mitigated without understanding the data that underpins model behavior.

The bill correctly treats data practices as inseparable from AI risk, but it does so almost entirely from the perspective of developers, deployers, and regulators. What is largely missing is a clear articulation of individual data rights. Opaque data practices undermine trust, accountability, and competitiveness not only because they impair oversight, but because they deny individuals meaningful agency over how data about them is collected, used, and reused. By contrast, risk-based data protection regimes such as the GDPR explicitly recognize that individual rights to access, correction, and transparency are not merely consumer protections, but mechanisms

²⁵ Rhea et al., “Resume Format, LinkedIn URLs and Other Unexpected Influences on AI Personality Prediction in Hiring,” *AAAI/ACM Conference on AI, Ethics, and Society*, 2022, <https://dl.acm.org/doi/10.1145/3514094.3534189>

for improving data quality and system performance. S.1169-A gestures toward data sensitivity and documentation, but stops short of recognizing individuals as active participants in data governance.

This gap matters. AI systems do not fail in isolation; they fail because of data choices made upstream and reinforced downstream, often without the knowledge or ability of affected individuals to identify errors or contest misuse. Treating data governance as fully participatory infrastructure would strengthen the bill's core objectives by surfacing errors earlier, improving data quality over time, and aligning incentives for responsible data stewardship across the AI lifecycle. While S.1169-A represents an important step forward, future amendments or implementing guidance should consider how individual data rights can be more directly supported, particularly in high-risk contexts where inaccurate or misused data can have lasting consequences for workers, consumers, and communities.

Why this strengthens competitiveness. Clear rules around data collection, use, documentation, and governance reduce both legal uncertainty and reputational risk for businesses. Companies that operate in opaque data environments face heightened exposure to enforcement actions, litigation, and public backlash, particularly as expectations around data protection and AI accountability continue to evolve. From a market perspective, predictability supports investment and innovation. Investors and enterprise customers consistently favor environments where data governance obligations are clear and stable, rather than subject to ad hoc enforcement or post hoc correction.²⁶ Firms that build AI systems designed to operate reliably within regulated environments gain a durable competitive advantage: they are better positioned to scale, to enter regulated sectors such as finance, healthcare, and employment, and to weather shifts in public expectations and regulatory scrutiny. In this sense, strong data governance does not constrain AI markets; it underwrites their long-term viability.

Background from my prior work. As I emphasized in my recent testimony before the Committee on Science, Space and Technology of the U.S. House of Representatives²⁷, opaque data practices undermine trust, accountability, and oversight across the AI lifecycle. Without visibility into data sources, preprocessing steps, and reuse practices, neither regulators nor affected individuals can meaningfully assess whether a system is safe, lawful, or fit for purpose. In my academic work, I have argued that effective AI governance must begin with data

²⁶ OECD, *Data Governance and Privacy in the Digital Age*, <https://www.oecd.org/digital/data-governance>

²⁷ Stoyanovich, "Deep Seek: A Deep Dive," Testimony at the Hearing of the Committee on Science, Space and Technology of the U.S. House of Representatives, Research and Technology Subcommittee, April 2025, <https://science.house.gov/2025/4/deepseek-a-deep-dive>

governance. The principle to “follow the data” is a prerequisite for functional AI oversight, because data choices shape model outputs, error modes, and downstream impacts long before a system is deployed. Absent strong data documentation and controls, even well-intentioned risk management efforts are likely to fail.

Comparison to the RAISE Act. The RAISE Act focuses primarily on model security, safety protocols, and incident reporting for frontier models. While these measures address important upstream risks, they do not directly confront the downstream data harms experienced by consumers and workers in everyday AI deployments. S.1169-A fills this gap by requiring developers and deployers to document and manage data practices in high-risk systems used across the economy. By centering data protection and data quality alongside model assessment, the bill provides a more complete and practical foundation for responsible AI markets in New York State.

Recommendation 5: Pair the Bill with Investments in Education, Auditor Expertise, and Government Capacity

S.1169-A assigns a central role to independent auditors and specifies requirements to ensure their independence (§ 87), and it relies on ongoing risk management, reporting, and enforcement to make its protections meaningful (§§ 89, 89-c). These provisions are well designed, but their effectiveness ultimately depends on the expertise, resources, and institutional capacity of those charged with implementing them.

The bill’s success will depend not only on what is written in statute, but on who implements it and how. Audits, risk management programs, and enforcement actions are only as strong as the technical and institutional capacity behind them. Capacity-building should therefore be treated as an essential component of enforcement, not as an optional or downstream concern. Without sufficient expertise among auditors, regulators, and enforcement staff, even carefully drafted requirements risk becoming superficial, inconsistently applied, or ineffective in practice.

Why this strengthens competitiveness. From an economic perspective, consistent and expert enforcement is a competitive asset. Companies are more willing to invest and innovate in jurisdictions where rules are applied predictably and competently, rather than sporadically or arbitrarily. Clear expectations, paired with knowledgeable oversight, reduce uncertainty and compliance whiplash. At the same time, regulation of this kind creates new markets for auditing, evaluation, documentation, and compliance services. New York is well positioned to become a hub for these activities, leveraging its existing strengths in technology, finance, and professional services. Investing in capacity thus supports both effective governance and economic development.

Background from my prior work. My experience with the implementation of New York City’s Local Law 144 on automated employment decision tools illustrates the risks of under-resourced oversight. While the law established important guardrails, limited auditor expertise and uneven regulatory capacity created confusion for employers, vendors, and affected individuals alike. I have spent years teaching regulators, practitioners, and members of the public about responsible AI precisely to address this gap. These efforts consistently demonstrate that education and training are prerequisites for meaningful accountability. Without them, even well-intentioned regulation can fall short of its goals.

Comparison to the RAISE Act. The RAISE Act acknowledges the importance of education and expertise, but does not specify how capacity should be built or sustained. S.1169-A presents an opportunity to do better at the state level by explicitly pairing substantive obligations with investments in auditor training, regulatory expertise, and public-sector AI literacy. Doing so would help ensure that the bill’s protections are enforced consistently and credibly, and that New York leads not only in setting standards for responsible AI, but in building the institutional capacity needed to uphold them.

Summary Statement

The New York AI Act (S.1169-A) is a timely and necessary step to protect New Yorkers from the adverse impacts of high-risk AI systems while strengthening the state’s long-term competitiveness. The bill’s risk-based scope, notice and opt-out protections, and requirements for auditing, reporting, and risk management provide essential market guardrails. To fully realize these goals, implementation should more clearly specify how validity and reliability are assessed, and New York should pair the bill with investments in government capacity and auditor expertise. Over time, the state should also strengthen individual data rights so that data governance becomes fully participatory infrastructure in high-risk settings. With these elements in place, New York can lead in building AI markets that are not only innovative, but trustworthy, durable, and aligned with civil rights and economic opportunity.