Testimony of Harriet P. Pearson, Managing Principal, Axia Advisory LLC

Before the U.S. Senate Committee on Health, Education, Labor, and Pensions

Hearing on "AI's Potential to Support Patients, Workers, Children, and Families"

Thursday, October 9, 2025

Chairman Cassidy, Ranking Member Sanders, and distinguished members of the Committee, thank you for the opportunity to testify today. I appreciate your commitment to thoughtfully examining AI's implications across the areas under this Committee's jurisdiction.

INTRODUCTION

I appear before you today speaking only for myself and not on behalf of any organization with which I have been or am affiliated. My remarks draw on more than thirty years of work at the intersection of information technology, business, and law.

Earlier in my career, I served as IBM's first Chief Privacy Officer—the Fortune 500's first such appointment—for twelve years, building and leading a global privacy and data protection program. During my total of nineteen years at IBM, I also held executive leadership roles in human resources, legal, and public affairs. My HR experience included HR leadership of the company's hardware business as well as companywide responsibility for HR strategy, HR policy, employee relations, and performance management. When I moved back onto the company's legal team, I served as Assistant General Counsel responsible for cybersecurity in addition to privacy and data protection.

In 2012, I became a Partner at the law firm Hogan Lovells, in what is now known as the Data, Privacy & Cybersecurity practice. I advised hundreds of clients across many industries for over a decade before retiring from the firm to serve as Executive Deputy Superintendent and head of the Cybersecurity Division at the New York Department of Financial Services until June 2024. While at NY DFS, I led the rewrite of the state's comprehensive cybersecurity regulation for that sector. I now divide my time between professional and community pursuits: through my consultancy I advise clients on digital governance matters; and I own and operate two small businesses with my husband: a 100-seat restored historic theater and a local news publication in the Eastern Panhandle of West Virginia.

In 2025, I was recognized as one of the top 20 practitioners in the past 20 years globally by the *Financial Times* Innovative Lawyers program, and in prior years received the Vanguard Award, the highest honor in the privacy field, from the International Association of Privacy Professionals; recognition as a Cybersecurity and Privacy Trailblazer from the *National Journal*; and top ranking in Privacy and Data Security law by Chambers & Partners.

So, my perspective today reflects my longstanding involvement in addressing the implications of new information technologies. I've been fortunate to have had the chance to think hard about what's possible, and what's practical, in formulating new privacy, cybersecurity, and employment laws, regulations, and policy. I've also been able to experience firsthand the positive and negative aspects of how such policies get implemented and their impact. I am a firm believer in the effectiveness of market forces to generate solutions and value for all Americans; but I am also open to legislative and regulatory solutions when they are clearly needed and appropriately designed.

For example, in the early 2000s, I testified before Congress in support of the Genetic Information Nondiscrimination Act—as I recall, I was one of the only business representatives to advocate for its enactment. I championed the need for a targeted federal law change to prohibit employment discrimination based on genetic information—such as genetic information showing a predisposition to develop a disease later in life. Congress enacted the landmark GINA legislation in 2008.

Today, we are still very much in the early days of generative AI development and deployment. While there has been some good work to identify risks and develop key principles for AI legislation, now is not yet the time for new AI-specific employment regulation. Instead, a proactive legislative focus on skills and education will be much more productive and effective.

THE POTENTIAL AND RISKS OF AI

Artificial intelligence holds tremendous promise for improving how we work, learn, receive healthcare, and solve complex problems. We should approach this technology with appropriate optimism about its beneficial applications while remaining clear-eyed about the risks when AI is not designed and deployed responsibly.

AI's beneficial potential includes:

- Enhancing productivity—including for smaller businesses like mine—by automating routine tasks and freeing up humans to do higher-value work requiring judgment, creativity, and interpersonal skills
- Improving decision-making by surfacing relevant information quickly and identifying patterns humans might miss
- Expanding accessibility for people with disabilities through assistive technologies
- Advancing medical diagnosis, drug discovery, and personalized treatment
- Personalizing education to individual learning needs and styles
- Strengthening cybersecurity by identifying threats and vulnerabilities faster
- Increasing workplace safety through predictive maintenance and hazard detection

The risks when AI is deployed inappropriately include:

- Errors in outcomes due to poor-quality data sets or training
- Opacity in decision-making that makes it difficult to understand why particular outcomes occurred
- Counterproductive workplace practices such as excessive workplace surveillance
- Over-reliance on AI systems that lack adequate human oversight
- Security vulnerabilities in AI systems themselves
- Erosion of human skills and professional judgment

The key insight is that AI is a tool. Like previous transformative technologies, its impact depends on how we choose to design, deploy, and govern it. Of course, even if all AI solutions are deployed well, changes in how work gets done will affect the individuals who do the work: changing many jobs, creating new jobs, and eliminating other jobs.

LEARNING FROM PRIOR TECHNOLOGY WAVES

It is instructive to consider AI in the context of prior technology-driven innovation waves. Each new wave—from the internet to mobile computing to cloud services—has been accompanied by initial hype, predictions of immediate transformation, and calls for urgent regulation. The reality has consistently been more measured: meaningful adoption takes years, implementation proves more complex than anticipated, and many concerns work themselves out through market forces, evolving best practices, and application of existing legal frameworks.

Generative AI, which burst into public consciousness with ChatGPT's release in late 2022, follows this familiar pattern. While adoption has been rapid by historical standards, we remain in early stages. According to a September 2025 Anthropic report citing the U.S. Census Bureau's Business Trends and Outlook Survey, AI adoption among U.S. firms more than doubled from 3.7% in fall 2023 to 9.7% in early August 2025. [1] Impressive growth, yet the data indicate that a large majority of firms does not yet use AI in production processes. A 2025 McKinsey survey found that 78% of organizations use AI in at least one business function,[2] but drilling down reveals that most are still piloting applications rather than fully integrating AI across operations. These findings ring true to me, reflecting accurately what I have seen in the organizations with which I work.

[1] The Anthropic Economic Index Report (Sept. 15, 2025) at p. 31, available at https://assets.anthropic.com/m/218c82b858610fac/original/Economic-Index.pdf

[2] McKinsey, The state of AI: How organizations are rewiring to capture value at p. 14, available at https://www.mckinsey.com/~/media/mckinsey/business%20functions/quantumblack/our%20insights/the%20state%20of%20ai/2025/the-state-of-ai-how-organizations-are-rewiring-to-capture-value_final.pdf

This measured pace is appropriate, particularly in high-stakes sectors like healthcare, financial services, and employment where organizations must carefully balance potential benefits against legal, operational, and reputational risks.

HOW GENERATIVE AI WORKS WITH PEOPLE

To understand AI's workplace implications, it helps to consider how generative AI actually assists humans in getting work done. Accenture's testimony before another Senate committee in 2023 offered a useful taxonomy of five core ways that generative AI works with people:

> **1. As an always-on advisor**, putting new kinds of intelligence into human hands in areas ranging from sales enablement and human resources to medical and scientific research and corporate strategy.
>
> **2. As a creative partner**, offering new ways to reach and appeal to audiences, bringing speed and innovation to production design, visual identity, copy generation, and real-time personalized marketing.
>
> **3. As a software developer**, boosting productivity by automating code writing, debugging, documentation, and predicting and pre-empting problems.
>
> **4. As an automation driver**, especially for tasks that provide historic context, present next best actions, summarize information, or make intelligent predictions.
>
> **5. As an enterprise protector**, helping companies use AI to their advantage in governance and information security, including in Security Operations Centers to mitigate threats and identify vulnerabilities faster.

This framework highlights an important point: AI primarily augments human capabilities rather than replacing human workers entirely. The impact is on tasks and workflows, not wholesale elimination of jobs.

WORK DISPLACEMENT VS. WORKFORCE DISPLACEMENT

This distinction between work displacement and workforce displacement is crucial for policy discussions.

AI will certainly displace certain tasks and change how many jobs are performed. History suggests, however, that technological change typically transforms work rather than eliminating it entirely. The automation of manufacturing, the digitization of office work, and the computerization of countless processes have repeatedly demonstrated this pattern.

Tasks likely to be augmented or displaced by AI include:

- Data entry and routine document processing
- Basic customer service inquiries and responses

- Preliminary legal document review and research
- Routine financial analysis and report generation
- Initial resume screening in hiring processes
- Transcription, translation, and summarization
- Basic diagnostic support in healthcare settings
- Routine coding and software testing

New work and roles being created include:

- AI trainers who teach AI systems to perform specific tasks effectively
- AI ethicists and governance specialists who ensure responsible deployment
- Prompt engineers who optimize human-AI interaction
- AI assessors and auditors who test systems for accuracy and compliance
- Data labelers and curators who prepare high-quality training data
- Human-in-the-loop specialists who review AI-generated outputs
- AI literacy trainers who help workers use these tools effectively
- Change management professionals who help organizations adapt workflows
- AI system integrators who connect AI tools with existing systems
- Explainability specialists who make AI decision-making transparent

The net employment effect remains uncertain and will vary significantly by industry, occupation, and geography. What seems clear is that many jobs will change substantially, requiring workers to acquire new skills and adapt to new ways of working. This reality places a premium on education, training, and lifelong learning—areas squarely within this Committee's expertise and jurisdiction.

WE ARE IN THE EARLY DAYS

Based on my advisory work with organizations across sectors and on publicly available data, we remain in the early stages of AI-enabled innovation in the workplace. While headlines suggest rapid and dramatic transformation, the truth on the ground is more nuanced.

What I see is that most larger businesses and organizations recognize AI's potential and are working on adopting at least some elements of it. However, as I previously noted, progress is appropriately measured, especially in higher-risk uses and sectors. In healthcare, financial services, legal services, employment, and other contexts, organizations typically proceed cautiously because the stakes are high and any potential benefits must be carefully balanced against multiple types of risk:

Legal risk: Potential violations of civil rights laws, privacy statutes, labor laws, or contractual obligations

Operational risk: System failures, inaccurate outputs, integration challenges with existing processes

Reputational risk: Loss of customer, employee, or public trust from poorly implemented AI

Compliance risk: Navigating an evolving patchwork of state and international AI regulations

Cybersecurity risk: Vulnerabilities in AI systems themselves or in the data they process

Implementing AI into significant workflows is a multi-step, involved process requiring substantial investment, technical expertise, careful planning, organizational change management, and ongoing monitoring. This does not happen overnight. Responsible organizations pilot applications, test for bias and accuracy, train workers, adjust processes, monitor results, and iterate—all of which takes considerable time and resources.

This reality counsels patience in policymaking. We should allow these processes to unfold, learn from what works and what doesn't, and avoid locking in regulatory requirements based on today's technology and today's limited understanding of long-term impacts.

WORKFORCE PRIVACY: A DISTINCT CONTEXT

I noted previously that one of the risks of inappropriately deployed AI systems can be excessive workplace surveillance. Given my background in privacy, I would like to address this more deeply, first noting that privacy in the employment context differs from consumer privacy in important ways that should inform policy discussions.

Legitimate Business Interests: Employers have recognized needs to monitor work performance, protect company assets and trade secrets, ensure workplace safety, comply with legal obligations including regulatory requirements, prevent harassment and misconduct, and manage operations effectively. Cybersecurity risk management in particular often necessitates some degree of user activity monitoring to detect insider threats, prevent data breaches, and protect both the organization and its employees from harm. In a 2021 report I co-authored examining workplace monitoring practices across 15 countries, we found that monitoring user interactions with data and systems has become foundational to cyber risk management programs, particularly as workforces have become more distributed and remote.[3] While these interests should be balanced against worker privacy, they are legitimate.

Reduced Privacy Expectations: Courts in the United States have long held that employees have diminished privacy expectations when using employer-provided equipment, working on employer premises, and during work hours. This legal reality reflects the employment relationship's nature. Privacy rights do not disappear in the workplace, but the balance differs from consumer contexts.

Power Asymmetry: Unlike consumer relationships, employment involves significant power imbalance. Workers cannot simply switch employers the way consumers switch products. They depend on employment for livelihoods and often for health insurance. This asymmetry counsels for thoughtful protections, but not necessarily protections identical to consumer privacy rights.

Existing Legal Frameworks: Multiple federal and state laws already address workplace privacy, surveillance, and data practices. These include Title VII and other civil rights statutes, the Americans with Disabilities Act, the Genetic Information Nondiscrimination Act, the National Labor Relations Act, the Electronic Communications Privacy Act, the Fair Credit Reporting Act, state constitutional privacy provisions, and various state laws on electronic monitoring, biometric data, and workplace surveillance. As our 2021 research demonstrated, organizations can successfully navigate these varied legal frameworks to implement necessary security measures while respecting employee privacy through careful program design, transparent policies, and appropriate limitations on monitoring scope.[4] These frameworks are technology-neutral and continue to apply to AI-powered systems.

EXISTING FEDERAL LAWS APPLY TO AI IN EMPLOYMENT

Multiple existing federal laws already provide substantial protections against potential harms from AI deployment in employment contexts:

Labor Law: The National Labor Relations Act (NLRA) prohibits employers from interfering with, restraining, or coercing employees' exercise of Section 7 rights. This prohibition is technology neutral.

Anti-discrimination Laws: Title VII of the Civil Rights Act prohibits discrimination in employment based on race, color, religion, national origin, or sex—whether such discrimination is intentional or results from facially neutral practices that have disparate impact. Employers are

[3] Harriet Pearson and W. James Denvil, *Protecting the Workforce and Information in a Global Landscape: A Legal Review* (Hogan Lovells, 2021), available at https://www.forcepoint.com/sites/default/files/resources/reports/report-hogan-lovells-protecting-the-workforce-and-information-in-a-global-landscape-en_0_0_0.pdf

[4] Id.

also prohibited from unlawfully discriminating based on age under the Age Discrimination in Employment Act and based on disability under the Americans with Disabilities Act.

Genetic Information: The Genetic Information Nondiscrimination Act (GINA) prohibits employment discrimination based on genetic information and restricts employers from requesting, requiring, or purchasing genetic information about employees or their family members. These protections apply regardless of how genetic information might be used—whether by human managers or AI systems.

Credit Reporting: The Fair Credit Reporting Act (FCRA) regulates the use of consumer reports in employment decisions and requires employers to provide notice and obtain consent before obtaining such reports.

Moreover, the vast majority of employers will not want to deploy AI-enabled tools in ways that undermine trust and efficacy in the workplace or that are contrary to industry-standard practices. Maintaining employee trust, attracting and retaining talent, and preserving organizational culture are powerful incentives for responsible AI deployment. Competition for talent, concern for organizational reputation, and potential legal liability all create strong incentives for employers to engage in responsible AI deployment.

Multiple industry organizations have developed voluntary best practices and frameworks that employers can adopt, including the Future of Privacy Forum's Best Practices for AI and Workplace Assessment Technologies and the National Institute of Standards and Technology's AI Risk Management Framework. Market forces, reputational concerns, and potential liability under existing laws drive meaningful adoption of such standards.

A RISK-BASED APPROACH TO TARGETED REGULATION

There is growing consensus—including among industry stakeholders—that a risk-based approach to AI regulation is appropriate. Such an approach applies heightened scrutiny and requirements to higher-risk use cases while allowing lower-risk applications to proceed with lighter oversight.

The Algorithmic Accountability Act of 2025 (S. 2164) is a step in this direction, although there are significant issues that would need to be addressed for a bill like this to be practical and not impose undue complexity and burdens on the private sector. The bill would require impact assessments for activities that use automated decision systems to make “critical decisions"—defined to include decisions with legal, material, or similarly significant effects on consumers' lives in areas such as employment, housing, healthcare, education, and financial services.

While the bill's emphasis on transparency and risk-based regulation is commendable, there is still work to be done to refine foundational definitions and focus the scope of such legislation. For example, given the specialized nature of certain sectors, should such laws be industry-specialized, or should one federal agency have overall jurisdiction? Given the potential for rapid technological evolution, are today's definitions sufficiently scoped or will they cover more than should be included? Consider that virtually any information system could be characterized as "input" to human decision-making, but that does not make every system a proper subject of specialized regulation. The focus should be on systems that effectively determine outcomes, not systems that merely inform human judgment.

Given that AI use in employment remains fairly limited at this point, as I've described, I am open in principle to precise, targeted legislation or regulation that addresses genuinely high-risk use cases where existing legal frameworks demonstrably fall short. However, several important considerations should guide any such effort:

First, any new requirements should be proportionate to actual risks and avoid imposing compliance burdens that would discourage beneficial uses of AI or disproportionately burden smaller employers.

Second, legislation should build upon—rather than duplicate or conflict with—existing anti-discrimination laws and their enforcement mechanisms.

Third, requirements should be technology-neutral and focused on outcomes rather than prescribing specific technical solutions that may quickly become outdated.

Fourth, we should take learnings from state-level experimentation. States like Colorado, California, and Illinois are testing different approaches, and this diversity of approaches will help identify what works and what is impractical.

Finally, even targeted regulation should proceed carefully given how rapidly the technology and its applications are evolving. What appears high-risk today may become routine tomorrow, and vice versa.

In short, while I support the principle of risk-based regulation for genuinely high-risk applications, the key question is whether the current limited deployment of AI in employment warrants new federal legislation now, or whether existing laws combined with agency guidance and state-level experimentation provide adequate protection while we learn more about the technology's actual impacts.

RECOMMENDATIONS

Given where we are—early in AI adoption with much still uncertain about long-term impacts—I recommend to the Committee a measured, principles-based approach:

1. Leverage Existing Frameworks First

Chairman Cassidy's prior work in this area wisely emphasizes adapting current frameworks before creating new ones. I strongly concur with this approach. Existing employment laws, civil rights statutes, and privacy regulations already provide substantial protections against the most serious potential harms from workplace AI.

Regulatory and enforcement agencies should be encouraged to look at how existing laws apply to AI-powered employment decisions. Courts will develop common law principles as cases arise. State legislatures are experimenting with various approaches. We should allow these existing mechanisms to work and learn from these varied approaches before assuming wholesale new federal legislation is required.

2. Use a Technology-Neutral Approach

Critically, any legislative approach to workplace issues should be technology-neutral rather than focused specifically on AI. Employment law should address discrimination, privacy violations, and unfair practices regardless of whether the employer uses AI, traditional software, or purely human decision-making. Technology-specific regulation risks becoming quickly outdated and may inadvertently favor certain technical approaches over others.

3. Prioritize Education and Workforce Development

Rather than extensive new regulation, substantial investment in education and workforce transition would be far more valuable:

AI Literacy: Help workers understand how to effectively use AI tools, recognize their limitations, spot potential biases, evaluate AI-generated outputs critically, and protect their privacy online and in the workplace. This should be integrated into K-12 education, community college programs, and workforce development initiatives.

Employer Best Practices: Support development and dissemination of voluntary best practices, model policies, and decision frameworks that help employers deploy AI responsibly. Industry associations and professional organizations are well-positioned to develop sector-specific guidance.

Workforce Transition Programs: Expand and modernize training programs to help workers acquire skills for an AI-enabled economy. This should include:

- Modernizing the Workforce Innovation and Opportunity Act to support reskilling and upskilling
- Expanding Pell Grant eligibility for high-quality short-term training programs
- Scaling apprenticeship and earn-while-you-learn models
- Supporting community colleges in developing AI-related curriculum
- Providing robust career counseling and transition services

4. Enact Comprehensive Federal Consumer Privacy Legislation

A useful step Congress could take is enacting comprehensive federal consumer privacy legislation. Such legislation would strengthen baseline privacy protections for all Americans across contexts, including some workplace scenarios involving consumer technologies (health apps, wearables, personal devices).

Everyone is a consumer—including every worker. Universal consumer privacy protections would benefit workers substantially without the complications of carving out special employment rules that might prove problematic given the legitimate differences between employee and consumer relationships.

5. Support Research and Data Collection

Congress should support research and data collection to better understand AI's actual impacts on employment, wages, working conditions, and workforce demographics. The Bureau of Labor Statistics and other agencies should track trends in AI adoption, job displacement and creation, wage effects, and skill requirements. This empirical foundation will inform whether and what legislative action becomes necessary.

6. Ensure Agency Oversight

Ensure that agencies charged with interpreting, updating, and enforcing existing employment laws—the EEOC, NLRB, Department of Labor, and others—are monitoring AI developments with implications for their areas of responsibility. Confirm that they have or are developing workplans to apply existing legal frameworks, resources, and technical expertise to AI-related issues in their scope.

7. Targeted Interventions Only Where Clearly Necessary

If specific problems emerge that existing frameworks demonstrably cannot address, Congress can make targeted adjustments. A risk-based approach that focuses regulatory requirements on high-risk applications—such as AI systems used in a material way to make consequential employment decisions—has merit and reflects growing consensus among stakeholders. However, any such legislation would need careful refinement to be practical and avoid imposing undue complexity and burdens. Such interventions should be:

- **Technology-neutral**: Focused on outcomes and harms rather than specific technologies
- **Evidence-based**: Supported by clear evidence of harm that existing laws cannot remedy
- **Proportionate**: Calibrated to actual risks rather than hypothetical worst-case scenarios
- **Flexible**: Able to adapt as technology evolves
- **Practical**: Recognizing compliance realities for employers of different sizes and sectors

CONCLUSION

Chairman Cassidy, your caution against one-size-fits-all AI regulation and your emphasis on adapting existing frameworks before creating new ones reflects sound judgment. I strongly support this measured approach.

We remain in the early days of AI in the workplace. The technology continues to evolve rapidly. Employers are still learning how to deploy it effectively and responsibly. Workers are adapting to new tools and ways of working. Many fundamental questions about AI's long-term impacts remain unanswered.

This is precisely the wrong time for sweeping federal legislation specifically focused on artificial intelligence in employment. As I testified two decades ago regarding genetic information, targeted legislation can be appropriate, but today's situation differs fundamentally. Such broad AI-specific legislation risks:

- Locking in outdated assumptions about technology that is still evolving
- Stifling beneficial innovation that could improve productivity and working conditions
- Creating rigid requirements that cannot adapt to technological change
- Preempting valuable state-level experimentation
- Imposing compliance burdens without commensurate benefits

Instead, this Committee can provide valuable leadership by:

- **Supporting education and training** to help workers thrive in an AI-enabled economy
- **Encouraging research and data collection** to ground future policy in evidence
- **Promoting voluntary best practices and industry standards** developed by those closest to the technology
- **Providing oversight of how** agencies are applying existing laws to AI systems
- **Supporting enactment of comprehensive federal consumer privacy legislation** that provides baseline protections across many contexts including workplace scenarios
- **Monitoring developments** to identify if and when targeted, technology-neutral interventions become necessary

The future of work is not predetermined. Through thoughtful policy that emphasizes education over regulation, evidence over speculation, and existing frameworks over new mandates, we can help ensure that AI serves as a tool for broadly shared prosperity.

I testified twenty years ago in support of GINA because genetic discrimination was a problem that existing laws did not adequately address, and the legislative solution at hand was targeted and practical. Congress significantly addressed the issue for the American people by simply adding "genetic information" to the list of protected characteristics that must not be used discriminatorily in employment. The law change was technology-neutral with minimal regulatory complexity. Today's situation with generative AI is fundamentally different. The technology is still emerging, its workplace applications remain limited, impacts are still unclear, existing laws already provide substantial protections, and where targeted interventions may eventually be warranted for genuinely high-risk applications, a carefully crafted risk-based approach would be more appropriate than broad technology-specific regulation.

I appreciate the opportunity to share these views and welcome your questions.

Harriet Pearson
HPP@AxiaAdvisoryLLC.com