

To: U.S. Department of Labor, Office of the Secretary

From: Chijioke Mgbahurike; Kwame Ocran

Date: 21 February 2025

Subject: Establishing a Federal Framework for AI Governance in Employment

Executive Summary

AI systems are increasingly used to screen applicants, rank candidates, evaluate performance, and support termination decisions. While these tools promise efficiency, they also scale risks that are difficult for workers and applicants to detect or contest, including discrimination, opaque decision logic, and error at high volume. Documented cases and reporting suggest that biased training data and proxy variables can produce unequal outcomes even when employers intend neutrality, and that automated systems can deny individuals meaningful explanation or recourse.

This memo proposes creating a federal regulator, the **Automated Employment Bureau (AEB)**, housed within the Department of Labor to establish baseline standards for the use of AI in employment decisions. The AEB would require tool registration and independent audits, mandate clear disclosures when automated tools materially shape employment outcomes and provide a formal complaint and redress pathway for applicants and workers.

AEB would reduce uneven protections created by state and local patchwork laws, strengthen [civil-rights](#) and [labor protections](#) in the context of algorithmic opacity, and provide employers and vendors with consistent national expectations. By pairing enforceable standards with periodic updating, this approach supports responsible innovation while ensuring that high-stakes employment systems remain transparent, contestable, and fair.

Policy Recommendation

The Department of Labor should support legislation to establish an Automated Employment Bureau (AEB) to regulate AI systems used in employment decisions.

Specifically, this memo recommends that the Department:

- Create the AEB within the Department of Labor with authority to set national standards for AI systems used in hiring, promotion, evaluation, and termination.
- Require registration of covered tools and independent audits before deployment and on a recurring schedule.
- Require standardized disclosures to applicants and workers when AI tools materially influence employment decisions, including meaningful explanation at an appropriate level of detail.
- Establish a complaint, investigation, and redress process for applicants and workers affected by unjust AI-based decisions.
- Enforce compliance through civil penalties and federal contracting leverage, with heightened oversight for repeat violations.

Background and Rationale

AI is no longer limited to experimental HR tools. Automated screening and ranking systems are now integrated into routine employment workflows, often delivered through third-party vendors and embedded in applicant tracking systems. This shift changes the scale of decision-making. A flawed rule, biased feature, or poorly validated model can affect thousands of people before anyone identifies the pattern.

Existing civil-rights and labor protections remain essential, but their effectiveness depends on transparency, documentation, and the ability to investigate decision processes. Many AI tools operate as opaque systems that are difficult for workers to challenge and difficult for regulators to assess without technical capacity. At the same time, state and local governments have begun to regulate automated employment decision tools through audits, disclosures, and appeal rights. Those efforts demonstrate viable policy instruments, but their geographic patchwork creates uneven protections and inconsistent requirements for employers operating across jurisdictions.

A federal regulator dedicated to AI in employment would extend familiar principles of nondiscrimination, accountability, and due process into the context of algorithmic decision systems. It would also create a mechanism for updating standards as tools evolve.

Problem Statement

The rapid deployment of AI tools in employment decisions creates three recurring public-interest risks that existing enforcement systems are not well-positioned to address at scale.

First, AI systems can reproduce and amplify discrimination. Models trained on historical employment data can encode patterns of unequal opportunity and replicate them through automated screening. Even when protected attributes are not explicitly included, proxy variables such as educational history, employment gaps, location, or language patterns can drive disparate outcomes. When employers cannot detect or correct these effects, applicants and workers face significant barriers to demonstrating discrimination because the decision logic and training data are often inaccessible.

Second, AI-driven systems create transparency and due process gaps. Many tools function as black boxes to affected individuals. Workers and applicants may not know whether an automated tool was used, what factors drove an outcome, or how to challenge errors. The absence of meaningful explanation makes it difficult to correct inaccurate information, contest flawed scoring, or identify systemic bias.

Third, automation increases the scale and speed of workforce disruption. Employers can use AI not only to screen candidates but also to support workforce reductions and reorganization. Workers with limited access to retraining, including older workers and those in low-income regions, face higher risk of displacement and less ability to adapt. These dynamics can interact with algorithmic bias, especially when displaced workers re-enter labor markets governed by automated screening.

Together, these risks justify a unified federal approach that establishes baseline standards, requires independent validation, and provides enforceable pathways for transparency and redress.

Current Policy Landscape

Current approaches to regulating AI in employment include collective bargaining, state and local regulation, and federal agency guidance. Union negotiations have produced AI-related protections in certain sectors, but low union density limits coverage and unions may lack the technical capacity to evaluate complex models consistently.

State and local governments have pursued more direct regulation. [Colorado's Senate Bill 24-205](#) establishes documentation requirements and appeal rights for AI-based employment decisions. [Illinois' House Bill 3773](#) prohibits AI-driven discrimination, including indirect discrimination through proxies, and sets transparency expectations. [New York City's Local Law 144](#) requires bias audits for automated employment decision tools and mandates disclosures. These models reflect important policy tools, including third-party auditing, transparency requirements, and contestability.

However, the current landscape remains fragmented. Multi-jurisdiction employers must comply with inconsistent definitions, audit standards, and disclosure requirements. Enforcement capacity varies across jurisdictions, and workers in states without AI-specific rules may have no comparable protections. The result is uneven worker protection, compliance complexity, and a high likelihood that harmful systems become entrenched before affected individuals can seek relief.

A federal baseline would preserve the strengths of existing state and local approaches while providing consistent standards and enforcement nationwide.

Proposed Policy

Congress should establish the Automated Employment Bureau (AEB) within the Department of Labor to regulate AI systems used in hiring, promotion, performance evaluation, and termination decisions.

AEB would have four core functions:

Standards and certification. AEB would define baseline national standards for fairness, transparency, and nondiscrimination for covered tools. It would require registration of covered systems and establish a pre-deployment certification process for tools that materially influence employment decisions.

Independent audits and ongoing monitoring. AEB would require independent audits before deployment and on a recurring schedule. Audits would assess disparate impact, validity, documentation practices, and performance compared to existing decision processes. AEB would

accredit third-party auditors and specify minimum audit content requirements to ensure consistency.

Transparency and notice. AEB would require employers to disclose when AI tools materially influence employment decisions and provide meaningful explanations of decision factors at a level appropriate for affected individuals. AEB would also maintain a public registry of certified tools and publish redacted audit summaries to support accountability while protecting legitimate proprietary information.

Redress and enforcement. AEB would operate a complaint intake and investigation process for applicants and workers. When violations are found, remedies could include corrective action plans, tool modification, changes to employer processes, reinstatement or back pay where appropriate, and mandated training. AEB would enforce compliance through civil penalties and federal contracting leverage, with heightened oversight for repeat violators.

Regulatory precedent for a dedicated bureau. Congress has created specialized federal regulators when existing agencies lacked the focused authority and technical capacity to address fast-evolving, high-impact markets. [The Consumer Financial Protection Bureau](#) is a clear example. It was established in Title X of the [Dodd-Frank Act](#) to set standards, supervise covered entities, and enforce consumer protections across a complex financial ecosystem. AEB would follow a similar logic for AI-driven employment systems by combining technical oversight, standard-setting, enforcement, and public accountability within a single bureau.

Implementation Plan

AEB should be established through enabling legislation that defines covered tools, sets enforcement authority, and provides stable funding. Implementation should proceed in phases to reduce burden and allow technical standards to mature.

Year 1: Authorizing legislation and initial rulemaking. Congress would pass an Automated Employment Bureau Act establishing authority, definitions, and enforcement mechanisms. The Department of Labor would initiate stakeholder consultation with employers, vendors, civil-rights organizations, labor groups, and technical experts.

Year 2: Stand-up and pilot certification. AEB would hire core staff with technical and legal expertise and publish initial technical guidance aligned with widely used risk and fairness frameworks. AEB would launch a pilot certification program focusing on a limited set of tools, such as systems used by federal contractors, to refine metrics and audit procedures.

Year 3: Certification and audit rollout. AEB would open registration broadly and require certification for covered tools by a defined compliance deadline. AEB would begin routine audits and publish aggregated results and enforcement statistics in an annual report.

Ongoing: Updating standards and public reporting. AEB would update standards on a regular schedule through notice-and-comment processes, maintain a public registry, and publish annual enforcement and outcomes reporting.

Stakeholders and Risks

This policy directly affects job seekers and workers, particularly those in protected classes who face elevated risk from biased screening and those with limited ability to contest automated outcomes. Employers and recruiters would benefit from clearer national expectations and reduced compliance fragmentation, though they would incur audit and reporting responsibilities. HR technology vendors would face stronger accountability but would gain market clarity through a consistent certification standard.

Potential concerns include regulatory burden, impact on innovation, and institutional overlap with existing agencies such as the Equal Employment Opportunity Commission and the Federal Trade Commission. These concerns can be mitigated by designing AEB as a specialized regulator with technical capacity and clear coordination mechanisms. AEB would not replace existing civil-rights enforcement; it would provide the technical oversight, standardized auditing infrastructure, and transparency mechanisms needed to make those protections workable in the context of algorithmic decision-making.

Conclusion

AI-driven employment decision tools can increase efficiency, but they also scale discrimination, error, and opacity in ways that existing worker protections struggle to address without transparency and technical oversight. State and local laws have begun to demonstrate workable audit and disclosure models, but the current patchwork produces uneven protections and inconsistent standards.

Establishing an Automated Employment Bureau within the Department of Labor would create a coherent federal framework that sets baseline standards, requires independent audits, mandates meaningful transparency, and provides enforceable pathways for redress. This approach protects applicants and workers while supporting responsible innovation and restoring accountability in high-stakes employment decisions.

