

AI-DRIVEN RECRUITMENT IN INDIA: LEGAL ACCOUNTABILITY AND THE FUTURE OF FAIR HIRING

Dr. Preeti Goel*
Niharika Aggarwal**

Abstract

Artificial intelligence (AI) is increasingly embedded in Indian recruitment through résumé screening, candidate ranking, chatbot interaction, digital assessments, video-interview analytics, and predictive matching. These systems promise speed and administrative efficiency, but they can also convert historical inequality into apparently neutral scores and recommendations. This paper examines how Indian law can allocate responsibility when AI-assisted recruitment compromises equality, privacy, accessibility, or procedural fairness. It argues that India's current legal position is fragmented: the Digital Personal Data Protection Act, 2023 (DPDP Act) regulates digital personal-data processing; disability legislation protects equality and reasonable accommodation; and constitutional privacy and equality norms provide interpretive foundations. Yet none creates a complete regime for automated employment decisions, algorithmic impact assessment, bias auditing, explanation, or contestation. The paper proposes an approach in which employers bear a non-delegable duty to validate, monitor, and review tools they deploy, while vendors carry enforceable duties of documentation, testing, security, and cooperation. The proposed approach combines risk assessment, independent bias audits, candidate notice, meaningful human oversight, accessible alternatives, audit trails, shared liability, and effective grievance redress. This would allow India to realise the productivity benefits of AI without permitting opaque systems to become instruments of exclusion.

Keywords: Artificial intelligence, recruitment, algorithmic bias, data protection, employment discrimination, legal accountability, fair hiring.

Introduction

Recruitment is a gateway to economic opportunity. A decision to shortlist, reject, rank, or interview an applicant affects access to livelihood, dignity, professional mobility, and social participation. Indian employers are

* Assistant Professor, Maharaja Agrasen Institute of Management Studies, GGSIPU, Delhi

** Student, BALLB, Maharaja Agrasen Institute of Management Studies, GGSIPU, Delhi

increasingly using digital tools to manage high volumes of applications, identify candidates, parse résumés, administer assessments, schedule interviews, and communicate with applicants. When these tools use machine-learning models or other AI techniques to infer suitability, rank applicants, predict performance, or automate workflow decisions, recruitment becomes an exercise in algorithmic governance rather than merely an administrative process.

The attraction is understandable. AI may reduce repetitive work, bring structure to large applicant pools, identify skills that conventional credential filters overlook, and help recruiters respond more rapidly. Yet the same system can produce unequal results at scale. A model trained on historic hiring records may learn that characteristics associated with previously selected workers are markers of merit. If past selection reflected gendered occupational segregation, language privilege, caste-linked educational access, ableist assessment practices, regional inequality, or preference for uninterrupted careers, the model may reproduce those patterns even without explicitly using protected attributes. Proxy variables such as postal code, college, employment gap, language style, device type, or social-media activity can operate as indirect channels for exclusion.¹ The principal challenge is accountability. Employers may contend that a vendor designed the system, vendors may characterise outputs as recommendations rather than decisions, and recruiters may claim that a human remained nominally “in the loop.” These explanations do not resolve whether applicants were treated fairly, whether their data were lawfully processed, or who must remedy a discriminatory result. The problem is intensified by opacity: an applicant may never know that AI was used, which data influenced the score, or how to seek meaningful reconsideration. This paper argues that AI-assisted recruitment should be treated as a high-impact employment practice. India should not permit an employer to evade responsibility merely because a third party supplied the tool or because a human formally clicked “approve” after an automated ranking. The legal framework must combine privacy, equality, accessibility, due process, and institutional accountability. The paper further explains the operation and risks of AI recruitment and analyses India’s present legal foundations. The paper identifies accountability gaps and proposes lessons and offers reforms for fair hiring.

1 Solon Barocas & Andrew D. Selbst, *Big Data’s Disparate Impact*, 104 *Calif. L. Rev.* 671, 692–700 (2016).

AI Recruitment and its Risks

A. From Automation to Algorithmic Decision-Making

Not all recruitment technology raises the same concern. A calendar tool that schedules interviews is different from a system that scores speech patterns, filters résumés, recommends candidates, or predicts “culture fit.” The legal significance increases when technology materially influences who sees an advertisement, advances to the next stage, receives an interview, or is rejected. Recruitment AI may include natural-language processing to extract résumé information, matching engines that compare profiles against job descriptions, automated tests, game-based assessments, chatbot screening, asynchronous video-interview tools, and predictive models that infer job performance or likelihood of retention. A common misconception is that automation necessarily increases objectivity. Algorithms make decisions through data, objectives, and design choices selected by people. A model’s apparent neutrality may conceal subjective assumptions about what counts as success, which variables are relevant, whose data are representative, and what level of error is tolerable. Technical accuracy alone is not legal or ethical fairness. A model may accurately replicate historical hiring patterns while still perpetuating unjust exclusion.

B. Bias, Proxy Discrimination, and Unequal Impact

Bias can emerge at every stage of the AI lifecycle. Training data may reflect historically unequal hiring; labels such as “high performer” may embody managerial prejudice; feature selection can rely on proxies for social advantage; and performance testing may fail to include candidates from different linguistic, geographic, gender, disability, or socioeconomic backgrounds. A résumé model trained on successful employees in a male-dominated technical workforce, for example, could rate indicators associated with male career patterns more highly. It might penalise maternity-related career breaks, part-time work, non-English communication styles, or educational pathways common among applicants with fewer resources.

India’s social context makes this problem especially important. Formal attributes such as institution attended, urban location, fluency in a particular language, continuity of employment, and access to high-speed internet are not distributed evenly. These features can become proxies for caste, class, rural origin, disability, or gendered care responsibilities. An AI system need not collect caste or religion explicitly to generate disparate outcomes. The legal inquiry should therefore assess effects, not merely the list of variables entered into the model.

C. Privacy, Surveillance, and Data Quality

AI recruitment frequently relies on more data than conventional recruitment. Applicants may provide résumés, identity documents, contact details, educational history, test results, video recordings, voice samples, behavioural metadata, and inferred personality traits.² Some systems may analyse facial movement, word choice, response time, or background information. This expansion creates risks of excessive collection, use for unrelated purposes, insecure sharing with vendors, retention beyond necessity, and decisions based on inaccurate or context-poor data. Privacy is not simply a matter of secrecy. In the employment context, data practices can shape opportunity. An inaccurate inference about communication style, emotional state, disability, or reliability can become a barrier to employment. Therefore, lawful recruitment requires data minimisation, purpose limitation, security, accuracy, and meaningful opportunities to challenge consequential errors.

D. Opacity and Procedural Unfairness

An applicant cannot effectively contest a decision that is invisible.³ A generic rejection email does not reveal whether a human assessor reviewed the application, whether the candidate was rejected because of a score, or whether a system incorrectly parsed the résumé. Trade-secret arguments cannot justify total secrecy where a person's employment opportunity is at stake. Fairness does not necessarily require disclosure of source code; it does require intelligible notice, explanation of the principal factors and process, access to correction, and a genuine avenue for human reconsideration.

India's Existing Legal Framework

A. Constitutional Values: Equality, Dignity, and Privacy

The best place to start is Articles 14, 16 and 21. Together they give us a way to judge whether AI hiring tools are fair, equal and respect privacy. Article 14 promises equality before the law⁴. In *E.P. Royappa v State of Tamil Nadu*, the Supreme Court said arbitrary state action goes against this and *Maneka Gandhi v. Union of India* added that when the State effects someones rights, it has to act fairly and reasonably.⁴ So if a government body uses an AI tool to shortlist candidates and the result turns out arbitrary, Article 14 can be used to challenge it. The catch is that Article 14 mostly binds the State, so

2 *Digital Personal Data Protection Act, No. 22 of 2023, Sec. 2(i), 2(k), 8, INDIA CODE (2023).*

3 *Andrew D. Selbst & Solon Barocas, The Intuitive Appeal of Explainable Machines, 87 Fordham L. Rev. 1085, 1087-1100 (2018).*

4 *Maneka Gandhi v. Union of India, (1978) 1 S.C.C. 248, 283-84 (India).*

it doesn't work the same way in a dispute with a private employer. It has the most force when a public authority or State instrumentality is doing the hiring.

Article 16 guarantees equal opportunity in public employment,⁵ and *Indra Sawhney v. Union of India* built up the law on fairness and non-discrimination there. This means that if a public sector recruiter brings in AI for screening, the same rules on equality and non-arbitrariness still apply. Article 16 doesn't directly cover private hiring though, so there's a real gap when a private company's AI system ends up discriminating. Article 21 protects life and personal liberty, and after *Maneka Gandhi* it also demands fair procedure whenever a protected right is affected. Then in *K.S. Puttaswamy (Retd.) v. Union of India*, the Court recognised privacy as a fundamental right.⁶ This matters for AI recruitment because these tools often handle applicant details, behavioural signals and digital profiles. If that data is collected or used without proper transparency and safeguards, it raises informational privacy problems.

B. Digital Personal Data Protection Act, 2023

For digitally processed applicant information, the DPDP Act is the most important general statute.⁷ It tries to balance an individual's right to protect their personal data against the need to process data for lawful purposes. In recruitment, an employer, agency or platform can be a data fiduciary,⁸ and a technology provider handling data for it can be a data processor. The Act's rules on notice, consent, legitimate use, security, accuracy, grievance redress and processor governance therefore apply directly to AI-assisted hiring. Several practical obligations follow for recruitment systems. Applicants should get a clear notice of what data is processed and why. Where data is used to make decisions about a person, it should be accurate, complete and consistent. Employers must put reasonable security safeguards in place and offer an accessible way to raise grievances. Vendor arrangements shouldn't become a way to dodge these duties: a company that picks an AI recruiter stays answerable for its data practices even when a cloud provider or HR-tech vendor does the technical work. That said, the DPDP Act 2023 is not an AI-employment statute. It doesn't expressly give anyone a right not to be subject to purely automated recruitment decisions, a right to a detailed explanation of an algorithm, or a mandate for bias testing. Data protection can control inputs and

5 *Indra Sawhney v. Union of India*, 1992 Supp. (3) S.C.C. 217, 419–20 (India).

6 *Justice K.S. Puttaswamy (Retd.) v. Union of India*, (2017) 10 S.C.R. 569, 638–39 (India).

7 *Digital Personal Data Protection Act*, No. 22 of 2023, Sec. 8, INDIA CODE (2023).

8 *Rights of Persons with Disabilities Act*, No. 49 of 2016, Sec. 20, INDIA CODE (2016).

procedures, but it can't tell us whether a model that looks lawful is still producing discriminatory hiring outcomes. That gap is why separate rules for employment AI are needed.

C. Disability, Accessibility, and Reasonable Accommodation

The Rights of Persons with Disabilities Act, 2016 is an essential lens for AI recruitment. Section 20 prohibits disability discrimination⁹ in employment matters in Government establishments and requires reasonable accommodation and a barrier-free, conducive environment. The section is limited by the kind of institution it covers, but the Act's principles of equality and accommodation should still guide how all recruitment is designed. Exclusion can happen with no discriminatory intent at all. A hiring portal that doesn't work with assistive technology, an automated video assessment that penalises atypical speech or eye contact, and a timed test with no accommodation are all examples. Accessible AI hiring should offer alternative assessment formats, keyboard navigation, screen-reader compatibility, flexible timing and accessible communication. Candidates should also be able to request accommodation without being penalised for it. Employers should test systems with disabled users before rolling them out instead of waiting for individual complaints. Accessibility has to be built in from the start, since fixing it after deployment comes too late.

D. The Fragmentation Problem

So Indian law has some useful pieces of protection but no single rulebook for employment AI.¹⁰ Privacy law covers data, wage law covers sex discrimination in recruitment, and disability law covers equality and accommodation. Constitutional doctrine supplies the underlying values, and general contract, negligence and consumer-law principles may help in particular disputes. What the framework doesn't clearly require is an algorithmic impact assessment, an external bias audit, a notice of adverse decisions, auditable documentation, or a defined split of liability between employer and vendor. That leaves candidates, firms, vendors, regulators and courts uncertain about where they stand.

9 *Rights of Persons with Disabilities Act, No. 49 of 2016, Sec. 20, INDIA CODE (2016).*

10 *Digital Personal Data Protection Act, No. 22 of 2023; Rights of Persons with Disabilities Act, No. 49 of 2016, Sec. 20, INDIA CODE (2016).*

Legal Accountability for AI Hiring

A. Employer Responsibility Must Be Non-Delegable

The employer should bear primary responsibility¹¹ for outcomes produced by a recruitment tool it chooses to deploy. The employer defines the vacancy, selects the vendor, determines whether a score is advisory or decisive, sets thresholds, controls the recruitment workflow, and receives the business benefit of automation. These facts make the employer the most appropriate first point of accountability for candidates. Outsourcing technology may allocate commercial risk between the parties, but it should not eliminate the employer's duty to provide a lawful and fair recruitment process. A non-delegable duty would require employers to conduct vendor due diligence, ensure data-protection compliance, assess discriminatory impact, provide accommodations, maintain human review, and operate a grievance channel. It would not mean that employers must understand every technical detail of a model. Rather, employers must obtain sufficient documentation and assurance to make a responsible deployment decision. An employer that cannot explain what a tool does, what data it uses, how it was validated, and how errors can be corrected should not rely on that tool for high-impact selection.

B. Vendor Duties and Shared Liability

Vendors cannot be treated merely as neutral suppliers.¹² They design or configure models, determine data architecture, make claims about accuracy, select validation approaches, and control significant technical information. Vendors should therefore owe duties of reasonable care and transparency. At minimum, they should provide documentation on intended purpose, known limitations, training and validation methodology, performance metrics, data categories, security controls, version changes, and appropriate human-use conditions. They should also facilitate independent audits and notify clients of material risks or identified defects. Shared liability is especially important when harm results from both design and deployment. For example, a vendor may supply a résumé-ranking model trained on unrepresentative historical data, while an employer may set an aggressive rejection threshold and fail to monitor gender outcomes. Neither actor alone explains the harm. Law and contract should permit responsibility to be allocated according to control and fault, while preserving the applicant's ability to seek a remedy from the employer that made the hiring decision.

11 *Digital Personal Data Protection Act, No. 22 of 2023, Sec. 8, INDIA CODE (2023).*

12 *Nat'l Inst. of Standards & Tech., Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1 (2023).*

C. Meaningful Human Oversight

A human being in the workflow is not enough if that person simply rubber-stamps an AI recommendation. Meaningful oversight requires authority, competence, and time.¹³ The reviewer must be able to override the output, understand material limitations, examine relevant context, recognise potential bias, and record reasons for the final decision. A recruiter should not be required to review thousands of algorithmic rejections individually; rather, the system should be designed so that high-impact adverse decisions receive proportionate scrutiny, particularly where a candidate seeks review or raises an accommodation issue.

D. Transparency, Contestability, and Evidence

Fair recruitment requires a minimum package of procedural rights.¹⁴ First, applicants should be told when AI materially influences screening, ranking, assessment, or rejection. Second, they should receive a concise and intelligible account of the role of the system, the categories of data used, and the existence of human review. Third, they should be able to correct inaccurate data, request reasonable accommodation, and seek review by a qualified human decision-maker. Fourth, organisations should keep audit logs and decision records sufficient to investigate complaints. These measures also solve an evidentiary problem. Algorithmic discrimination is often systemic rather than intentional. Requiring an applicant to prove the inner workings of a proprietary model is unrealistic. A burden-shifting approach is preferable: once an applicant presents credible facts suggesting discriminatory or procedurally unfair treatment, the employer should demonstrate that the system was appropriately assessed, validated, monitored, and used with meaningful oversight.

Lessons for Fair and Accountable AI Hiring in India

A. Treat Recruitment AI as High-Impact Decision-Making

India's approach should begin with the recognition that AI tools used to recruit, rank, filter, assess, or reject applicants are not ordinary back-office technologies. They mediate access to employment, income, occupational mobility, and social dignity. Their effects can be particularly consequential where an automated score determines whether an applicant proceeds to an interview, sees a job advertisement, or is excluded before any human recruiter examines the application. Regulatory obligations

13 Nat'l Inst. of Standards & Tech., *Artificial Intelligence Risk Management Framework* (AI RMF 1.0), NIST AI 100-1 (2023).

14 Andrew D. Selbst & Solon Barocas, *The Intuitive Appeal of Explainable Machines*, 87 *Fordham L. Rev.* 1085, 1087-1100 (2018).

should therefore follow the impact of the decision, not the vendor's description of the product as a mere "assistive" or "efficiency" tool. A risk-based classification would allow Indian regulation to distinguish low-risk administrative automation, such as interview scheduling, from systems that make or substantially influence consequential judgments. Candidate ranking, automated rejection, predictive scoring, targeted vacancy advertising, behavioural or biometric analysis, and systems that infer personality or "fit" should fall within a high-impact category. This approach is more proportionate than a blanket ban on recruitment technology while recognising that technical convenience cannot justify opaque exclusion.

B. Build a Rights-Based Indian Framework

The Indian framework should be anchored in equality, dignity, privacy, accessibility, and procedural fairness. Articles 14, 15, 16, and 21 of the Constitution establish¹⁵ the normative basis for non-arbitrary and rights-respecting recruitment, particularly in public employment. The Supreme Court's recognition of privacy as intrinsic to life and personal liberty reinforces the principle that personal information cannot be collected, profiled, and used without regard to legality, necessity, proportionality, and safeguards. These constitutional values should inform both public-sector recruitment and the regulation of private hiring systems. They are especially relevant in India, where data points that appear facially neutral such as educational institution, postal code, language proficiency, employment gaps, device use, or location may function as proxies for class, caste, gender, disability, regional background, or unequal access to opportunity. A framework concerned only with whether an algorithm explicitly records a protected attribute will be inadequate. The legal inquiry must also examine whether the system's design and outcomes create unjustified disparate effects. The Digital Personal Data Protection Act, 2023 supplies important but incomplete protection.¹⁶ The Act provides a legal basis for regulating digital personal-data processing and gives data principals rights to correction, erasure, and grievance redress. It also requires reasonable security safeguards and imposes obligations on a data fiduciary in relation to personal data under its control. In recruitment, those duties should translate into clear applicant notices, accurate and relevant data, secure vendor arrangements, limited retention, and practical procedures for correction and complaint. However, data protection alone cannot resolve algorithmic discrimination. The DPDP Act

15 INDIA CONST. arts. 14, 15, 16, 21.

16 Digital Personal Data Protection Act, No. 22 of 2023, Sec. 8, 12-14, INDIA CODE (2023).

does not expressly create a comprehensive right against solely automated hiring decisions, require employers to disclose model logic, or impose regular fairness audits. A dedicated employment-AI framework should complement the Act by creating process-specific duties for organisations that use AI to determine access to work.

C. Make Employers Primarily Accountable

The employer should remain the primary point of accountability to an applicant. It is the employer that defines the role, procures or approves the tool, decides how heavily to rely on the output, sets decision thresholds, and receives the operational benefit of automation. These facts justify a non-delegable duty to ensure that an AI-assisted recruitment process is lawful, fair, accessible, and subject to meaningful review. An employer should not be able to answer a discrimination or privacy complaint simply by saying that the system was supplied by a third-party vendor. Primary employer responsibility should include: pre-deployment due diligence; validation of the tool for the specific job and applicant population; privacy and security review; accessibility testing; monitoring of adverse outcomes; human-review arrangements; candidate communications; and operation of an effective grievance channel. Employers need not become software engineers, but they must obtain enough documentation and assurance to assess whether a system is suitable for a high-impact employment use. If an organisation cannot identify the tool's purpose, data categories, known limitations, validation results, and route for correcting errors, it should not use that tool as a basis for adverse selection. This allocation of responsibility is consistent with existing Indian employment and data-protection values. The duty should apply with equal force where discrimination is generated through an algorithmic ranking or screening process rather than an individual recruiter's overt preference.

D. Impose Corresponding Duties on Vendors

Employer accountability should not make vendors invisible. Vendors design or configure models, determine the architecture through which information is processed, set validation methods, and often possess the technical information required to identify error, bias, and security risks. They should therefore have enforceable duties to document intended use, disclose material limitations, maintain version histories, support auditability, apply appropriate security practices, and inform clients of material changes or known defects. The appropriate model is shared but differentiated liability. Employers should answer to candidates for the fairness of the recruitment process; vendors should be responsible for defects in design, negligent validation, misleading claims, inadequate

documentation, or failure to provide promised safeguards. Contracts should allocate operational responsibilities but should not prevent injured applicants from seeking a remedy from the employer that relied on the AI output. Public procurement templates and sectoral guidance can make these duties concrete through clauses on audit access, data use, deletion, incident notification, sub-processors, model updates, and cooperation in investigations.

E. Require Algorithmic Impact Assessments and Independent Audits

Before deploying high-impact recruitment AI,¹⁷ organisations should conduct an Algorithmic Impact Assessment (AIA). The assessment should state the employment objective; explain why AI is necessary; identify the data used; evaluate privacy, discrimination, accessibility, and cybersecurity risks; consider less intrusive alternatives; specify the role of human oversight; and identify mitigation and redress measures. An AIA must be a living document, revised when the model, data source, job role, scoring threshold, or deployment setting changes. Periodic bias audits should accompany the AIA. Audits should test selection rates, false positives, false negatives, calibration, error patterns, proxy variables, and accessibility outcomes across relevant groups. In the Indian context, testing should not assume that a narrow binary-gender or urban-English-speaking dataset is representative. Depending on lawful data availability and the sensitivity of the context, organisations should examine whether system outcomes disadvantage women, persons with disabilities, linguistic minorities, rural candidates, applicants from economically weaker backgrounds, or persons with non-linear career histories. The objective is not to require unlawful profiling; it is to identify and mitigate foreseeable disparate effects.

F. Ensure Transparency, Human Review, and Contestability

Candidates should receive timely and intelligible notice when AI materially influences recruitment. The notice should identify the stage at which AI is used, the principal categories of information processed, whether a human will review the output, and the available route to seek correction, accommodation, and reconsideration. Meaningful transparency does not require publication of source code or disclosure of trade secrets. It requires enough information for an applicant to understand that an automated process has affected the opportunity and to challenge an error or unfair outcome. Human oversight must be substantive rather than ceremonial. A reviewer should have authority to override an automated recommendation, access to relevant contextual information, appropriate

17 Nat'l Inst. of Standards & Tech., *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1 (2023).

training, and sufficient time to conduct a real assessment. Accordingly, high-impact recruitment systems should provide alternative assessment formats,¹⁸ screen-reader compatibility, keyboard navigation, captioning, flexible timing where appropriate, multilingual information, accessible support, and a straightforward process for accommodation requests. Employers should test systems with persons with disabilities and diverse users before deployment. Inclusion also demands caution regarding tools that purport to infer emotion, personality, truthfulness, or employability from facial expression, voice, or behavioural signals. Unless such tools can satisfy strong standards of necessity, scientific validity, reliability, and non-discrimination, they should not be used to make adverse employment decisions.

H. Establish Coordinated Enforcement and Remedies

A fair-hiring framework will be ineffective without accessible enforcement.¹⁹ The Data Protection Board can address obligations within the DPDP Act, while labour authorities, disability-rights institutions, and sectoral regulators may have complementary roles. India should develop coordinated guidance or a specialised mechanism for high-impact employment AI. This mechanism could issue templates for AIAs, candidate notices, audit reports, procurement clauses, vendor contracts, and accessibility checklists; receive complaints; and publish anonymised compliance trends. Effective remedies should include data correction, explanation, human reconsideration, accommodation, suspension of unsafe tools, administrative sanctions for serious violations, and compensation where permitted by applicable law. Organisations should also preserve records of model versions, decision thresholds, audit findings, material vendor changes, and complaint outcomes. These measures promote accountability and give regulators a practical evidentiary basis for identifying systemic patterns rather than treating every case as an isolated error.

Conclusion

AI may help Indian employers manage recruitment efficiently, but efficiency is not a sufficient measure of legitimacy. Employment decisions shape life chances, and automated systems can amplify existing inequality at a speed and scale that traditional hiring practices cannot match. The principal legal question is not whether algorithms are inherently good or

18 *Rights of Persons with Disabilities Act, No. 49 of 2016, Sec. 20, INDIA CODE (2016).*

19 *Digital Personal Data Protection Act, No. 22 of 2023; Rights of Persons with Disabilities Act, No. 49 of 2016, Sec. 20, INDIA CODE (2016).*

bad. It is whether institutions deploying them can demonstrate that their systems are necessary, proportionate, fair, secure, accessible, and open to challenge.

India already possesses important legal building blocks. The DPDP Act establishes duties relevant to applicant data; disability law foregrounds equality and accommodation; and constitutional doctrine protects privacy, dignity, and equality. Yet these protections remain incomplete when an applicant is rejected by an opaque score or ranking model. The absence of an explicit framework for automated employment decisions leaves too much to voluntary ethics and private contractual arrangements. India need not choose between innovation and rights. It can encourage AI-enabled efficiency while insisting that recruitment technologies are tested, documented, accessible, and contestable. The central lesson is simple: a system that determines access to employment should be governed by standards commensurate with its social impact. The appropriate response is a high-impact employment-AI regime built around non-delegable employer responsibility, enforceable vendor duties, algorithmic impact assessments, independent bias audits, candidate notice, meaningful human review, accessible design, auditable documentation, and practical remedies. Such a regime would not prohibit innovation. It would direct innovation toward systems worthy of public trust. Fair hiring in India's digital future requires more than faster selection; it requires accountable selection.