Frontier software

Human Capital Management
& Payroll Software/Services

AI and the Pay Gap

The Legal Risks Employers Can't Ignore

Whitepaper

November | 2025

Imagine this: an organisation rolls out a sleek AI-driven performance tool designed to calculate bonuses. It pulls in data on productivity, collaboration, and goal completion, then delivers its verdict; a 15 per cent bonus for one employee and only 8 per cent for another, despite their records looking almost identical. When the results are challenged, the algorithm reveals patterns that seem to penalise certain demographic groups. Suddenly, a tool that promised efficiency looks suspiciously like it is entrenching bias. A frightening question then emerges: who is responsible when a machine makes a discriminatory pay decision?

This is not a futuristic thought experiment. It is a very real dilemma for Australian businesses already experimenting with algorithmic decision-making in performance management and compensation. The attraction is obvious; speed, consistency and the appearance of objectivity. The legal and ethical implications, however, are far murkier. Organisations must now navigate questions that straddle law, technology, and culture, often with little precedent to guide them.

Law Written For People, Not Machines

Anti-Discrimination Framework

Australia's anti-discrimination framework was written with human decision-makers in mind. Employers remain bound by the Racial Discrimination Act 1975, the Sex Discrimination Act 1984, the Age Discrimination Act 2004, the Disability Discrimination Act 1992, and the Fair Work Act 2009. However, none of these statutes make an exception for decisions made by software rather than people.

Intent vs. Outcome

The legal test becomes slippery when intent is considered. A manager can be questioned about the reasoning behind a decision, but how do you interrogate an algorithm whose outcomes emerge from millions of data points, weighted and processed in ways even the developers may not fully understand?

Compliance Gaps

Consider the Fair Work Act, which prohibits adverse action on discriminatory grounds. Determining intent is relatively straightforward for human managers. For AI, intent is less relevant than outcome, but the law has yet to catch up. Without clarity, employers risk compliance gaps, even when they are diligently trying to integrate technology responsibly.

Bias That Looks Like Neutrality

The danger of algorithmic bias lies in its subtlety. An AI system trained on years of historical performance reviews may replicate patterns reflecting past organisational prejudice. If historical promotions or bonuses favoured particular groups, the algorithm will likely reproduce those disparities. What makes this particularly insidious is the appearance of neutrality. The system applies the same rules consistently, so its outcomes appear objective, even while they are systematically unfair.

Over time, this repetition creates what some experts call "normalised discrimination." Employees see patterns emerge and assume that consistency equates to fairness. In reality, disadvantage is being hard coded into the system. For affected staff, recognising bias becomes difficult, as discrimination is cloaked in the authority of technology. Managers new to the organisation may never see the original discriminatory decisions; all they notice are consistent, "objective" outputs.

Subtle biases can manifest in unexpected ways. AI may weight communication style, work pattern preferences, or collaboration approaches that correlate with gender, age, or cultural background. A system that unintentionally penalises employees who prefer asynchronous communication or flexible hours may be inadvertently reinforcing inequality. Without oversight, these patterns can persist, embedding unfairness into the organisational fabric.

Whose Hands Are On The Wheel?

Australian courts have yet to tackle algorithmic employment discrimination in depth, leaving organisations navigating grey legal territory. Employers remain ultimately responsible for employment practices, but the introduction of AI raises complex questions: if an algorithm produces flawed outcomes, is liability solely with the employer, or does it extend to software vendors and developers? Should laws focus on intent when algorithms act independently of human instruction, or should they pivot and assess impact?

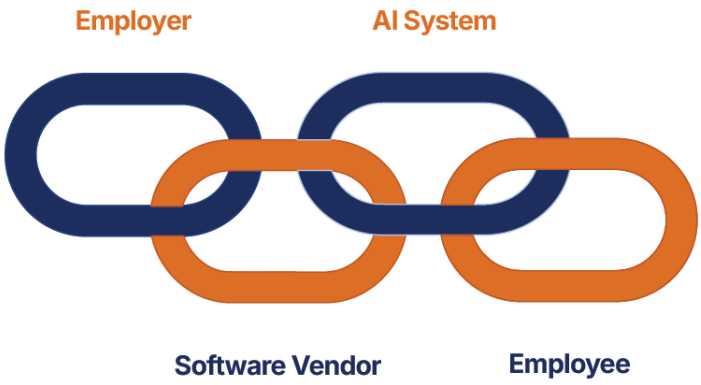

The complexity of modern AI compounds the challenge. Proprietary algorithms often operate as “black boxes,” leaving organisations unable to explain exactly how outcomes are determined. Yet transparency obligations remain. Leaders must reconcile commercial confidentiality with the need to provide staff with meaningful explanations.

For now, precedent suggests responsibility rests with employers. Even so, this clarity does little to ease the practical difficulties of explaining, monitoring, and controlling decisions that are increasingly driven by machines.

Culture and Trust on the Line

Legal compliance is one dimension, but the cultural stakes are even higher. Few employees will accept “the computer says no” as a credible explanation for why their bonus is lower than expected. Trust in pay systems is fundamental to workplace culture, and once that trust is undermined, disengagement and attrition often follow.

Transparency matters. Explaining how AI influences outcomes, reassuring staff that human oversight exists, and offering avenues to challenge decisions can make the difference between a workforce that embraces new tools and one that feels devalued by them. Leaders must view AI not as a replacement for human judgment, but as a tool whose output requires interpretation and contextualisation.

The human element remains central. Employees must see that algorithms are part of a system, not the final authority. Clearly communicating this helps preserve organisational credibility and ensures that innovation does not come at the expense of fairness.

Guarding Against Risk

Scrutinize Training Data

Examine historical bias in past bonus allocations, performance ratings, or promotion patterns before deploying algorithms.

Implement Human Oversight

Validate AI decisions, particularly those with large pay implications or patterns affecting specific demographic groups.

Document Thoroughly

Maintain clear records explaining how decisions are reached, the weighting factors, and the logic applied.

Manage Vendor Relationships

Ensure providers have rigorous bias testing, ongoing monitoring, and shared responsibility for discriminatory outcomes.

For organisations using or considering AI in pay, the path forward is proactive, not reactive. Historical bias left unchecked will only be magnified. Leaders should begin by scrutinising the data that trains these systems. Did past bonus allocations, performance ratings, or promotion patterns reflect disadvantages? If so, those patterns need correction before the algorithm is deployed.

Finally, grievance procedures provide a safety valve. Employees must have accessible, human-driven mechanisms to challenge AI-driven pay decisions. Without this, organisations risk eroding trust and engagement.

Australia And The World

Internationally, the debate around AI and fairness is accelerating. Europe's AI Act introduces requirements for transparency and auditing of high-risk systems, including those used in employment decisions. The Act, which came into force on 1 August 2024, applies directly across EU Member States and includes provisions that specifically address AI in recruitment, performance evaluation, and workforce monitoring. Breaches can result in fines of up to €35 million or 7% of global turnover for serious infringements.

In the United States, several states, including New York and California, have introduced or piloted legislation requiring algorithmic audits or disclosures for automated employment tools. These laws aim to mitigate harm, ensure transparency, and establish accountability in AI-driven hiring and compensation systems.

Canada is also moving toward a regulatory framework through its proposed Artificial Intelligence and Data Act (AIDA), which emphasises accountability, transparency, and human oversight in high-impact AI systems. While still in development, AIDA signals a clear intent to regulate workplace AI with a focus on fairness and risk management.

By comparison, Australia is only beginning to outline its regulatory approach, with early inquiries and commission reports signalling concern but offering limited concrete guidance.

The message for employers is clear: compliance expectations will rise. Those organisations that build transparency, accountability, and auditability into their AI processes now, will be better positioned when legislation catches up. Early preparation is not just prudent; it is strategic.

More Than Compliance

Ultimately, the algorithm question is about more than law. It is about fairness, trust, and human dignity. Technology offers efficiency and scale, but without careful governance it can entrench disadvantage under the guise of objectivity.

The organisations that succeed will treat AI as a tool to be managed responsibly, not an oracle to be obeyed. Leaders must ensure that AI enhances workplace fairness rather than hollowing it out. Machines may process the data, but people live with the consequences.

Approaching AI with intentional governance, regular review, and transparent communication allows leaders to harness its benefits while safeguarding employees. This is the frontier of modern remuneration, where technology and human judgment intersect, and where leadership shapes not just pay outcomes but organisational culture itself.

Transform Payroll into a Guardian of Fairness and Compliance
AI is reshaping the workplace – but without oversight, it can entrench bias and risk legal exposure. Frontier Software helps organisations turn payroll into a proactive force for equity, transparency, and regulatory resilience. Our solutions empower leaders to audit decisions, validate fairness, and build trust in a data-driven world. Visit www.frontiersoftware.com to learn more.

References

1. AKS Law. (2025, March 21). *AI in Recruitment: Legal Risks Australian Employers Need to Know in 2025*.
2. Australian Human Rights Commission. (2024). *Australia Needs AI Regulation*.
3. Australian Human Rights Commission. *Guidance Resource: AI and Discrimination in Insurance*.
4. Fair Work Ombudsman. *Protection from discrimination at work*.
5. HR Leader. (2025, February 12). *Update Fair Work Act to cover AI and automation, report says*.
6. Jackson, T. (2021). *Ethics and discrimination in artificial intelligence-enabled recruitment practices*. Humanities and Social Sciences Communications.
7. Lewis Silkin. (2023, October 31). *Discrimination and bias in AI recruitment: a case study*.
8. European Parliament and Council of the European Union, 2024. *Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence and amending various regulations and directives (Artificial Intelligence Act)*. Official Journal of the European Union, L series, 12 July.
9. Government of Canada, 2022. *Government of Canada introduces bill to strengthen privacy protections and create new framework for artificial intelligence*.

About Frontier Software

Frontier Software is a leading software and services organisation specialising in digital transformation for payroll and HR solutions. With over 40 years of industry experience, our comprehensive solutions are trusted by businesses across the globe. Our expert team is dedicated to helping organisations streamline their payroll and HR processes, ensuring compliance, accuracy, and efficiency. Looking to enhance your payroll and HR processes? We can help you optimise your existing Frontier Software solution or create a completely new system tailored to your business needs.

Frontier software

> **Commitment to customer service is our number one priority.**

Michael Howard, Founder

This document is provided for informational purposes only. While we have taken care to ensure the accuracy and reliability of the information presented, the content reflects our understanding at the time of publication. Circumstances, facts, and interpretations may change, and the material should not be regarded as advice or relied upon as the sole basis for making decisions. We accept no liability for any loss or damage arising from the use of this publication or from any inaccuracies, omissions, or changes that may occur after release. Readers are encouraged to seek independent advice or verification where necessary. Nothing in this document should be regarded as a commitment, warranty, or guarantee by the company. We are under no obligation to update or correct the information unless required by law.

Copyright © Frontier Software 2025