

FOR THE RECORD

The record is the defense.

AI moved into hiring. Federal enforcement moved out. States moved in. And the research on what actually predicts performance was rewritten. **One artifact sits at the center of all four shifts: the documented record of how a hiring decision was made.** This brief covers the law, the litigation, and the science, in plain language, with sources.

PREPARED FOR

HR & People Leaders

READING TIME

12 minutes

EXHIBITS

8

trugementalent.com

Every HR leader is being asked the same question.

Where does AI belong in hiring? Boards want the efficiency. Legal wants the exposure map. Candidates want an honest shot. Per the World Economic Forum, nearly nine in ten companies now use AI somewhere in candidate screening. The tools arrived faster than the rules. Then, over eighteen months, the ground moved three times.

SHIFT 01

Federal enforcement retreated.

The EEOC closed its docket of disparate impact charges and now routes those claims straight to court. The agency stepped back. The underlying liability did not move an inch.

SHIFT 02

States filled the gap.

Illinois, California, Colorado, New York City, Texas. Different definitions, different notices, different penalties. All aimed at the same thing: automated influence over hiring decisions.

SHIFT 03

The science was revised.

A 2022 re-analysis corrected decades of selection research. Structured interviews, human judgment against defined criteria, now sit at the top of the validity hierarchy.

~9 in 10

companies use AI somewhere in candidate screening (World Economic Forum)

0.5–2× salary

the cost of one bad hire on a critical role (SHRM; DOL floor is 30% of first-year earnings)

1 artifact

connects all three shifts: the record of how the decision was made

The referee left the field. The lawsuits didn't.

In 2025 the federal government pulled back from enforcing disparate impact, the theory that a neutral-looking practice can be unlawful if it disproportionately screens out a protected group. Many employers read that as lower risk. Read the fine print. The opposite is closer to the truth.

○ APR 23, 2025

Executive Order 14281 signed.

"Restoring Equality of Opportunity and Meritocracy" directs all federal agencies to deprioritize enforcement of disparate impact liability.

○ SEPT 2025

The EEOC closes its docket.

An internal directive instructs offices to discharge pending charges based solely on disparate impact.

○ OCT 2025

Right-to-sue letters go out.

Claimants are told the agency will not investigate. If they want to pursue the claim, the path is a private lawsuit in court.

○ DEC 10, 2025

DOJ rescinds Title VI regulations.

The Justice Department finalizes removal of disparate impact provisions from its own rules, and other agencies follow.

WHAT DID NOT CHANGE

● **The statute.**

Congress wrote disparate impact into Title VII in 1991. An executive order cannot repeal a statute. Courts still must apply it.

● **The private right of action.**

Individuals and the plaintiffs' bar can still sue. A right-to-sue letter skips the agency queue and puts the claim in front of a judge sooner.

● **State law.**

New York, California, Illinois, Colorado, and Minnesota all retain disparate impact liability in their own statutes and rules.

WHAT IT MEANS

The administrative filter is gone. Claims that once sat in an agency queue now move directly to a courtroom, and employment lawyers expect **more** private litigation, not less. In court there is no investigator to persuade. There is a record to produce, or a silence to explain.

One job posting. Multiple legal systems.

As Washington stepped back, the states stepped in. The rules differ on definitions, notice, and penalties. They agree on one thing: if technology influences a hiring decision, the employer answers for the outcome.

Illinois

EFF. JAN 1, 2026

HB 3773 · HUMAN RIGHTS ACT

AI-driven discrimination is a civil rights violation **even without intent**. Notice required when AI touches employment decisions. Zip codes banned as proxies for protected classes. Enforced with a **private right of action**.

California

EFF. OCT 1, 2025

FEHA · AUTOMATED-DECISION RULES

A discriminatory outcome from an automated system violates FEHA, intent or not. Employers must **retain ADS data four years**: inputs, outputs, criteria, bias-test results. Bias testing, or its absence, is evidence in a claim.

Colorado

EFF. JAN 1, 2027

SB 26-189 · REPLACED 2024 AI ACT

The landmark AI Act was repealed and replaced in May 2026 before it ever took effect. The lighter version still requires **disclosure and human review** for consequential decisions, including hiring.

New York City

EFF. 2023

LOCAL LAW 144

Automated employment decision tools require an **independent bias audit within the past year** before use, plus candidate notice. Applies to employers and employment agencies alike.

Texas

EFF. JAN 1, 2026

HB 149 · TRAIGA

Prohibits AI systems developed or used with **intent** to unlawfully discriminate. Narrower than Illinois or California: disparate impact alone does not establish intent under the Texas law.

The compounding problem

Post one remote-eligible role to candidates in New York, Illinois, and Colorado and you can trigger **three separate compliance regimes** with different definitions, notices, and penalties. The shared requirement across all of them: be able to show your work.

The case every vendor contract should be read against.

ACTIVE

Mobley v. Workday, Inc.

N.D. CAL. · No. 3:23-cv-00770 · IN DISCOVERY, MID-2026

THE FACTS ON FILE

100+

A rejected applicant alleges he applied to more than one hundred jobs through employers using Workday's screening tools and was rejected every time, sometimes within hours, including overnight. He claims the tools discriminated by **age, race, and disability**.

~1.1B

Applications processed through the platform during the relevant period, according to figures disclosed in the certification proceedings.

~14K

People reported to have opted into the certified age-discrimination collective by the March 2026 deadline. Workday declined to confirm the figure and disputes the claims.

WHY IT MATTERS TO EVERY EMPLOYER

LESSON 01 The vendor is not a shield.

The court allowed claims to proceed on an **agent theory**: employers plausibly delegated the function of advancing and rejecting candidates to the vendor's system. A tool acting on your behalf can carry your liability with it. Buying the tool outsources the work. It does not outsource the risk.

LESSON 02 A black box leaves you without a record.

When screening logic lives inside a vendor's model, you cannot produce your own account of why a candidate advanced or didn't. If the decision is ever challenged, that silence is the exposure. You are defending an outcome you cannot explain.

THE QUIET IMPLICATION

This is not an argument against AI tools. It is an argument about **where the record lives**. Whatever screens, ranks, or scores on your behalf, you carry the outcome. Carry a record of judgment with it.

Legal to destroy at month twelve. Needed at month thirty.

Federal rules require hiring records, including interview notes, to be kept for one year. But the moment a charge is filed, every related record must be preserved until the matter fully ends, through litigation and appeals. The gap between the legal minimum and the practical need is where organizations get hurt.



A contested hire rarely surfaces inside twelve months. Charges get filed, right-to-sue windows run, litigation stretches. By the time a question arrives, the notes are discarded, the spreadsheet is overwritten, and memory is filling in the blanks. **Memory is not evidence.**

1 year

The federal floor. EEOC rules (29 CFR 1602) require personnel and hiring records kept one year.

3+ years

The realistic window a contested hire stays live once a charge is filed: preservation runs until final disposition, including appeals.

4 years


California's retention rule for automated-decision data: inputs, outputs, criteria, and bias-test results.

The best predictor of performance is also the most defensible.

For twenty years, selection research said cognitive ability tests were the strongest predictor of job performance. In 2022, Sackett, Zhang, Berry, and Lievens re-examined the underlying data and corrected a systematic statistical error. The hierarchy changed.

WHAT PREDICTS JOB PERFORMANCE

Operational validity (r), Sackett et al., 2022, Journal of Applied Psychology



Scale shown to $r = .50$. The revised analysis places education and tenure-based credentials among the weakest predictors in the hierarchy; exact figures vary by measure, so no bar is drawn.

THE HONEST CAVEAT

The same study found **wide variance** around the structured-interview average. Structure done casually earns none of the advantage. The gain comes from defined criteria, independent scoring, and evidence behind every rating.

THE CONVERGENCE

Read that list again as a lawyer would. Defined criteria. Consistent process. Documented judgments. The practices that make a decision **accurate** and the practices that make it **defensible** are the same practices. Rigor and defensibility are not competing goals. They produce the same artifact.

AI can screen. It cannot decide. And it cannot defend.

The question is not whether to use AI in hiring. Most companies already do. The question is where the line sits. Draw it between processing information and exercising judgment, and the legal, scientific, and practical answers all line up.

● WHERE AUTOMATION EARNS ITS KEEP

- ✓ **Volume at the top of the funnel.** Sorting, scheduling, parsing, deduplicating. Machine work, done at machine scale.
- ✓ **Structuring information.** Normalizing applications, organizing evidence, preparing material for human review.
- ✓ **With the required notices.** Illinois and New York City already mandate disclosure. Build notice in from the start, not as a retrofit.

● WHERE JUDGMENT MUST STAY HUMAN

- ✗ **The decision on critical roles.** Plant managers, engineering leaders, executives. The hires where being wrong costs six figures.
- ✗ **Anything you may have to explain later.** A model's ranking is not a rationale. If it can be challenged, it needs a documented human reason.
- ✗ **The final call, anywhere.** Colorado's replacement law and California's rules both point the same direction: human review of consequential decisions.

INTERNAL TOOL OR THIRD PARTY: SAME MATH

Illinois's rules reach employers **and their agents**. California reaches anyone acting on an employer's behalf. Mobley tests whether a vendor is an agent. Whoever builds the tool, the deployer answers for the outcome. So apply one test before any tool touches a hiring decision: **if this hire is challenged in two years, can I produce a record of human judgment that explains it?** If the answer is no, the tool is in the wrong seat.

The anatomy of a defensible decision.

Strip away the statutes and the case law and a defensible hiring decision has four parts. They are the same four parts that make a decision good. Miss any one and the other three lose their force.

PILLAR 01

Criteria

Were the standards set before you met anyone?

Role-specific, weighted, agreed in advance. Criteria written after the fact read as justification, not judgment, and a reviewer can tell the difference. So can a court.

PILLAR 02

Independence

Did reviewers score before the room talked?

Independent judgments preserve the signal. Discussion-first panels converge on the most senior voice, and real disagreement, often the most valuable data in the room, gets averaged away.

PILLAR 03

Evidence

Is every score tied to something observed?

Impressions fade and cannot be cross-examined. A rating anchored to a specific answer, work sample, or behavior holds its value at month one and at month thirty.

PILLAR 04

Record

Is it complete, retrievable, and immutable?

A record that can be quietly edited later protects no one. Completeness matters too: missing data should be visible as missing, never dressed up as a safe-looking blank.

Five moves, in order.

1

Inventory every tool that touches hiring.

Internal or vendor. If it scores, ranks, filters, or recommends people, it is in scope, whatever the contract calls it.

2

Map your notice obligations.

Illinois and New York City already require disclosure. California's regime is live. Match every state you hire in against its rules, remote roles included.

3

Extend retention past the federal floor.

Twelve months is the minimum, not the standard. Keep complete decision records at least three years. Four if California touches your funnel.

4

Structure the decision itself.

Defined criteria, independent scoring, evidence behind every rating, documented rationale. For every critical hire, every time. This is also, per the research, simply how you pick better.

5

Test your own memory.

Pick a critical hire from eighteen months ago. Try to reconstruct how the decision was made: who scored, against what, and why the call went the way it did. **What you cannot produce is your exposure.**

Sources and further reading.

FEDERAL

Executive Order 14281, “Restoring Equality of Opportunity and Meritocracy” (Apr. 23, 2025). · **EEOC internal directive** on disparate impact charges, as reported by the Associated Press (Sept. 2025). · **DOJ final rule** rescinding Title VI disparate impact regulations (Dec. 10, 2025). · **29 CFR Part 1602**, EEOC recordkeeping requirements.

STATE & LOCAL

Illinois HB 3773 (amending the Human Rights Act) and Illinois Department of Human Rights draft rules. · **California** Civil Rights Council regulations on automated-decision systems under FEHA. · **Colorado** SB 24-205 (2024) and SB 26-189 (2026). · **New York City** Local Law 144 (2021). · **Texas** HB 149, the Responsible Artificial Intelligence Governance Act.

LITIGATION

Mobley v. Workday, Inc., No. 3:23-cv-00770 (N.D. Cal., filed 2023). Certification, agent-theory, and discovery rulings through mid-2026. Allegations are contested and no finding of liability has been made.

RESEARCH

Sackett, P. R., Zhang, C., Berry, C. M., & Lievens, F. (2022). Revisiting meta-analytic estimates of validity in personnel selection. *Journal of Applied Psychology*, 107(11). · **Society for Industrial and Organizational Psychology**, commentary on the revised hierarchy. · **SHRM** replacement-cost research (one-half to two times annual salary). · **U.S. Department of Labor** bad-hire cost floor (30% of first-year earnings). · **World Economic Forum**, AI adoption in candidate screening.

LEGAL ANALYSES CONSULTED

Jackson Lewis · Seyfarth Shaw · Duane Morris · Epstein Becker Green · Buchanan Ingersoll & Rooney · Perkins Coie · Ogletree Deakins · Maynard Nexsen · Akin Gump · Franczek P.C.

A NOTE ON METHOD

Every claim in this brief traces to a primary source or a named legal analysis, verified as of July 2026. Where a figure is reported but unconfirmed, the brief says so. Where the research carries a caveat, the caveat is printed. That standard is the point of the document.

ABOUT TRUGEM

Built for the decisions you may have to defend.

TruGem is the decision layer above your hiring process. Your panel scores candidates independently against defined, role-specific criteria. Real disagreement is surfaced instead of averaged away. Missing data looks like missing data, never like a safe score. And every hire produces a permanent, auditable Decision Report of how the call was made.

Humans decide. TruGem makes the decision structured, less biased, and defensible.

See how TruGem works → trugemtalent.com