



ACF IN ACTION

acfstandard.com

ACF en Action applique l'Agentic Commerce Framework® à des entreprises réelles, à partir de sources publiques. Chaque cas montre qui reste responsable quand l'IA décide.

ACF-CS-020 · Case Study #020 · Édition 2026

Mobley v. Workday (AI-driven screening of job applications)

Can a screening decision about people be delegated to an agent without measuring bias?

CARTE D'IDENTITÉ

SECTOR HR / software	COUNTRY USA	TECHNOLOGY Applicant screening	TYPE Agentic AI	AUTONOMY autonomous
CONFIDENCE High				

OUTILS ACF ILLUSTRÉS DANS CE CAS

ACF-01	ACF-02	ACF-03	ACF-08	ACF-00	ACF-05
--------	--------	--------	--------	--------	--------

mobilisé · à prévoir · 6 outils du Toolkit (4 mobilisés, 2 à prévoir, sur 17).

1. Executive summary

Workday provides many employers with an HR platform whose functionality relies on artificial intelligence to screen and recommend job applications: the system compares the skills required by a posting to candidate profiles and provides the employer with a recommendation. Starting in 2023, one candidate, Derek Mobley, then other plaintiffs aged over 40, brought a lawsuit. They **allege** that they were rejected almost systematically, without an interview, because of what they claim to be a bias in the recommendation system (age, but also race and disability invoked).

At this stage these are **allegations** brought before a court, not an established fault. Read through the lens of ACF, this case is the **"incident"** counterpart of the NYC Local Law 144 case, its reference: both address the same object, the **bias of a delegated decision**, along opposite paths. The case is instructive not as a condemnation, but as an illustration of the **risk** of delegating the screening of people to an agent without a fairness measure. The value of ACF here is **corrective**: an agent that screens people without a bias-measurement mechanism exposes the organization to a major risk, whether or not there is fault in the legal sense.

2. Context

Like many large employers, those who use Workday delegate part of the screening of applications to an automated system, in order to handle at scale volumes that human recruiters could not absorb. The intent is legitimate and common. The decision delegated to the agent this time concerns people: to set aside or advance job candidates, which engages equal treatment. The question of the case: can the screening of people be delegated without measuring the fairness of the outcome?

3. Public sources used

- Federal court (N.D. Cal.), Mobley v. Workday, Inc., 3:23-cv-00770 (public documents).
- Holland & Knight, "Federal Court Allows Collective Action Lawsuit Over Alleged AI Hiring Bias" (2025).
- Seyfarth Shaw, analysis of the ruling (the "agent" theory).

4. Reconstructed AI architecture (from public sources)

Workday application recommendation feature (at the heart of the dispute).

- Nature: a component of an HR platform relying on AI to compare the skills of a posting to candidate profiles.
- Function: to screen and recommend candidates to the employer, that is, in practice, to help set aside or advance profiles.
- Scope at issue: the measurement of bias (age, and other protected characteristics invoked) was, according to the plaintiffs, absent or insufficient.

The object of the dispute. Starting in 2023, Derek Mobley and other plaintiffs aged 40 and over **allege** that they applied to hundreds of jobs through employers using Workday and were rejected almost systematically, without an interview, due to a bias they attribute to the system. These facts are, at this stage, allegations, not established faults.

The state of the proceedings. In 2024, the court **allowed the disparate-impact discrimination claim** (under the ADEA and ADA anti-discrimination laws) to proceed, holding that Workday may be viewed as an **agent** of the employers who use its tool. In May 2025, the court granted **preliminary certification of a collective action** (for candidates aged 40 and over). The case has not been decided on the merits: the facts remain allegations.

Safeguards at issue (what the dispute questions):

- the existence, or not, of a **fairness measure** before and during deployment;
- the presence, or not, of an **indicator** of the selection rate by protected characteristic;
- the traceability of the measure, of the model version, and of the accountable party.

5. ACF mapping

 **Tools used:** ACF-01 (mapping and shared accountability) · ACF-03 (anti-bias prohibitions) · ACF-05 (fairness supervision) · ACF-08 (register).

In ACF language, the case reads as follows (without prejudging fault):

ACF element	In the Workday case (allegations)
Named human (accountable)	The employer who decides on the hiring and, according to the theory adopted by the court, the vendor as an agent of the employer. Accountability potentially shared.
Agent	The application recommendation system.
Platform	The Workday HR platform integrated into the employers' recruitment processes.
Delegated decision	To screen and recommend candidates, that is, to help set aside or advance people.
Autonomy level	High in practice on mass screening: many profiles set aside without an interview, according to the plaintiffs.
Non-delegable zones	The final decision to set aside a candidate: a decision about a person, of high criticality, to be kept under human control and to be measured.
Rule / safeguard	At the heart of the dispute: was a fairness measure in place? Without it, any potential bias remains invisible.
Control third party	After the fact, the court (Federal court, N.D. Cal.). No fairness measure published upstream is established.
Traceability	Open question: did the register trace the fairness measure, the model version, and the accountable party?

A schema in ACF graphic language accompanies this case (separate file).

6. Analysis

 **Tools used:** ACF-01 (shared accountability) · ACF-03 (anti-bias prohibitions) · ACF-05 (supervision) · ACF-08 (register).

Reminder: the facts complained of are allegations. The analysis bears on the governance risk, not on the guilt of a company.

What the case reveals (the governance risk):

- **Delegating screening does not erase accountability.** To believe that entrusting the screening to a tool transfers accountability to the vendor alone is a mistake. Accountability can be **shared**: the employer answers for its hiring decisions, and the court held that the vendor may be viewed as an agent of the employer.
- **A bias cannot be seen without measurement.** Without a fairness indicator, any potential bias remains invisible: the organization can neither observe it nor correct it. The risk lies in the **absence of measurement**, independently of any legal fault.
- **Removing age is not enough.** Proxies (dates, career paths) reintroduce the bias. It must be measured on the outcome, not merely by removing a variable.
- **Supervising performance is not supervising fairness.** Optimizing for the "good" candidate in the sense of the past can reproduce a bias. Fairness is a distinct supervision layer.

What ACF would have made explicit (and enforceable):

- an **agentic constitution** setting absolute anti-bias prohibitions: do not use variables correlated with age (or other protected characteristics), measure fairness before deployment, do not deploy as long as neutrality is not demonstrated;
- a **fairness indicator** tracked and published: the selection rate by age band (and other protected characteristics) compared to the qualified population;
- a **register** tracing the fairness measure, the model version, and the accountable party, so as to be able to answer for it.

7. ACF toolkit used

This grid indicates the tools that the ACF analysis uses (✓) to read this case. It does not mean the company uses ACF nor that a fault is established: it shows which instruments the framework would bring to measure and control the risk.

Tool	Status	Role in this case
ACF-01 · Decision map	✓	Map the shared accountability vendor / employer
ACF-03 · Agentic constitution	✓	Set the anti-bias prohibitions (correlated variables, no deployment without neutrality)
ACF-05 · Supervision	✓	Track and publish a fairness indicator
ACF-08 · Decisions register	✓	Trace the fairness measure, the model version, and the accountable party
ACF-02 · Criticality matrix	→ SOON	Classify the "set aside a candidate" decision as high criticality
ACF-12 · Agent mandate	→ SOON	Bound what the screening agent may recommend on its own

8. Governance questions (what an AI committee should ask itself)

1. Who answers for a screening decision delegated to a tool: the vendor, the employer, or both?
2. Can a bias be seen without a fairness measure?
3. Does the risk depend on an established fault, or on the absence of a measurement mechanism?
4. Is removing a protected variable enough, or must the bias be measured on the outcome (including via proxies)?
5. Does our supervision bear on performance, on fairness, or on both?

9. General lessons

A decision that bears on people cannot be delegated without being **measured**. The Workday case illustrates a risk, without prejudging fault: a screening of applications delegated to an agent without a fairness indicator exposes the organization, because any potential bias then remains invisible. And delegating the screening does not erase accountability for the hiring decision: it can be shared between the employer and the vendor.

The transferable lesson: a decision about people is measured and published, otherwise you discriminate without knowing it. The governance risk arises from the absence of measurement, even before any legal fault.

10. What ACF would do (reference design)

Prescriptive and non-critical formulation: here is how an organization would design such a system directly with the framework, whatever its sector.

If an organization wished to deploy this type of system directly with ACF, it could in particular:

- **map** accountability (ACF-01): identify that the employer answers for its hiring decisions and that the vendor may be an agent of the employer, therefore co-accountable;
- **declare** in a signed agentic constitution (ACF-03) absolute anti-bias prohibitions: no variables correlated with a protected characteristic, mandatory fairness measurement before deployment, no deployment without demonstrated neutrality;
- **keep** the final decision to set aside a candidate under human control, as a high-criticality zone;
- **track and publish** a fairness indicator (ACF-05): selection rate by protected characteristic, compared to the qualified population;
- **document** a decisions register (ACF-08) tracing the fairness measure, the model version, and the accountable party.

None of this is a judgment on Workday: it is the way ACF would make the fairness of a delegated decision that bears on people **measurable, signable, and auditable**.

11. Similar cases (transferability)

This case is specific neither to Workday nor to HR. It is transferable to any organization that delegates to an agent a **decision that bears on people and engages equal treatment**:

-  recruitment and talent management (CV screening, candidate scoring)
-  credit and insurance (scoring, pricing)
-  public services and allocation of benefits
-  healthcare (prioritization, triage)
-  education (admission, guidance)
-  justice and security (risk assessment)

Everywhere, the same pattern: an agent that screens people at scale, a decision (to set aside, refuse, prioritize) of high criticality, and a fairness measure without which a bias remains invisible. This is what makes this document a **reference case**: it illustrates a concept, not just a company.

12. Confidence of the assessment

High. The case rests on **public court documents** (Mobley v. Workday, Inc., 3:23-cv-00770, N.D. Cal.) and on analyses by leading law firms (Holland & Knight, Seyfarth Shaw). The state of the proceedings (the claim allowed to proceed in 2024, preliminary certification of a collective action in May 2025) is therefore established and verifiable by the reader. **Essential point: the facts complained of remain allegations, the case not having been decided on the merits.** The ACF analysis bears on the **governance risk** and presumes no fault.

Teacher guide

Pedagogical objective. Convey that a decision about people is measured and published, that delegating a screening does not erase accountability, and how to handle a case **without prejudging** fault (ongoing litigation, allegations).

Session outline (90 min).

1. (15 min) Reading of the **two** cases in the kit: Workday (incident, allegations) and NYC Local Law 144 (reference), without the analyses.
2. (20 min) In groups: map the shared accountability vendor / employer (card ACF-01).
3. (20 min) Workshop: write the anti-bias prohibitions (card ACF-03).
4. (20 min) Debate: "Can a bias be seen without a fairness measure?"
5. (15 min) Synthesis: a decision about people is measured and published, otherwise you discriminate without knowing it.

Important reminder to convey to students. This case concerns ongoing litigation. The facts complained of are allegations, not established faults. The analysis bears on the governance risk, not on the guilt of a company.

Possible exam questions. See the "Governance questions" section.

Reminder: grading students' work is a pedagogical act of the teacher. This guide provides the reference analysis, not automatic grading.

Associated lab (ACF-LAB-020)

From this case, the student produces:

- the mapping of the shared accountability vendor / employer (card ACF-01);
- three enforceable anti-bias prohibitions (card ACF-03): correlated variables, measurement before deployment, no deployment without neutrality;
- the fairness indicator to track and publish (card ACF-05): selection rate by protected characteristic;
- the schema of a decisions register (card ACF-08): fairness measure, model version, accountable party.

Appendix - Completed ACF cards

The official Toolkit cards used in this case, filled in with its public data. For illustration: an organization of the same profile can reuse and adapt them.



OBJECTIVE

Map a decision that bears on people and the accountability shared between the employer and the vendor. Filled from public sources, without prejudging fault (litigation ongoing).

1 AGENTIC MATURITY SCALE

0

Classic automation

Fixed rules, no autonomy.

1

Assisted agents

The agent proposes, the human decides.

2

Governed agents

The agent acts within a frame, under supervision.

TARGET ~80%

3

Supervised autonomy

The agent decides alone; control after the fact.

2 YOUR KEY DECISIONS (filled)

A

Match skills and profiles

STAKE

Compare a posting to applications

IMPACT

Aid to screening, under control

DEADLINE

Continuous

AUTONOMY

Suggestive (recommendation)

B

Recommend surfacing a profile

STAKE

Prioritize candidates to the employer

IMPACT

Fairness measure required

DEADLINE

Continuous

AUTONOMY

Co-decision + measure

C

Reject a person from hiring

STAKE

Screening decision on a person

IMPACT

Commits equal treatment

DEADLINE

-

AUTONOMY

NON-DELEGABLE

Reading: the decision to reject a person is high-criticality and stays under human control, with a fairness measure. Delegating screening does not erase accountability: it can be **shared** employer / vendor (agency theory retained by the court).



OBJECTIVE

Rate each decision by **criticality**, never by volume. On people, the irreversibility of a rejection and the equality stake place the final screening at the very top of the scale.

1 THE FORMULA

$$\text{Criticality} = (\text{Impact} \times \text{Frequency} \times \text{Irreversibility}) \div 10$$

Thresholds: < 15 **low** · 15-25 **moderate** · > 25 **high**

2 THE RATED DECISIONS (filled)

Type	I	F	Ir	Criticality	Zone / autonomy
Flagging an incomplete profile	2	6	1	1.2	Low · Autonomous
Skills / profile matching	3	9	2	5.4	Low · Suggestive
Prioritization recommendation	6	8	5	24	Moderate · Fairness measure
Mass screening without interview	8	8	6	38.4	High · Mandatory measure
Reject a candidate	9	7	7	44.1	High · Non-delegable

A decision that bears on people ranks as high-criticality. Without a fairness measure, a possible bias stays **invisible**: the risk arises from the absence of measurement, before any fault. The facts alleged against Workday remain allegations.



OBJECTIVE Define the agent's constitution: identity, standing principles and forbidden zones. Card filled with public case data.

1 AGENT IDENTITY & MISSION

AGENT NAME

Candidate recommendation system

ACCOUNTABLE (DDAO)

To be designated (employer; vendor as the employer's agent)

PRIMARY MISSION

Screen and recommend applications at scale

SCOPE

Aid to screening under a fairness measure; excluding the final decision to reject without human control

2 CORE PRINCIPLES

- 1 Measure fairness before and during deployment, by protected characteristic
- 2 Do not deploy until neutrality is demonstrated
- 3 Keep the final decision to reject a person under human control
- 4 Trace the fairness measure, the model version and the accountable owner

EXAMPLE PRINCIPLES

- Respect the legal and sector framework
- Protect data and confidentiality
- Stay transparent, every decision traceable
- When in doubt, escalate rather than decide

3 FORBIDDEN ZONES (ABSOLUTE)

- 1 Use variables correlated with age or a protected characteristic
- 2 Deploy a screening of people without a fairness indicator
- 3 Treat an automatic rejection as a final decision without human recourse

EXAMPLE PROHIBITIONS

- X Commit spending above a threshold without validation
- X Share personal data with a third party
- X Change its own rules or scope on its own

Constitution filled retrospectively, **without prejudging fault**. The facts alleged against Workday are allegations brought before a court, not established discrimination. Removing age is not enough: proxies reintroduce it, bias must be measured on the outcome.



OBJECTIVE

Record the fairness measure, the model version and the accountable owner, to be able to answer for a decision that bears on people. Filled with the fields relevant to the case.

1 FIELDS TO RECORD

Field to record	Example, Workday case (allegations)
Decision recorded	Recommendation to reject or surface a profile
Fairness measure	Selection rate by age bracket (open question)
Model version	Reference of the version that produced the recommendation
Accountable owner	Employer; vendor as the employer's agent
Status	Litigation ongoing, allegations not judged on the merits

2 USE OF THE REGISTER

- 1 Make a possible bias visible and measurable by protected bracket
- 2 Reconstruct which model version produced which recommendation
- 3 Establish who answers for the decision: employer and vendor, shared accountability
- 4 Provide enforceable evidence, without prejudging fault

Without a register tracing the fairness measure, a possible bias stays **invisible** and the organization can neither observe nor correct it. Reminder: the facts alleged are allegations, the case is not judged on the merits.

Disclaimer

This study is carried out **exclusively from public information** and concerns an **ongoing lawsuit**: the facts complained of are **allegations**, not established faults. It constitutes **neither an audit, nor a certification, nor legal advice**, nor an official evaluation of Workday. It applies the ACF framework for pedagogical and illustrative purposes. No internal, confidential, or non-public data was used. The trademarks cited belong to their respective owners.

Sources

- Federal court (N.D. Cal.), *Mobley v. Workday, Inc.*, 3:23-cv-00770 (public documents)
- Holland & Knight, *Federal Court Allows Collective Action Lawsuit Over Alleged AI Hiring Bias* (2025)
- Seyfarth Shaw, analysis of the ruling (the "agent" theory)

ACF Case Study ACF-CS-020 · v1 · working document, to be reviewed before any public distribution.