FOUR CONTROLS, ABOUT ONE ANALYST-DAY A QUARTER

Who is responsible for what the AI decided

In *Mobley v. Workday* a federal court held for the first time that an AI hiring vendor can be liable as an **agent** of the employer under Title VII, the ADEA and the ADA. You chose the tool and you handed the screening function to the vendor, but the legal obligation that goes with that function stayed with you.

Work down this sheet in order, tick each box, date it, and keep it on file. The record of having asked is itself part of your defence.

THE THREE GAPS, IN EVERY COMPANY NAMED

No adverse impact data asked for in the RFP

COST TO CLOSE: TWO TO FOUR HOURS

Left open, it produced a federal suit with defence costs estimated at $500k to $2M and rising.

No monitoring of rejection rates by group

COST TO CLOSE: ONE ANALYST-DAY A QUARTER

Left open, it ran unseen for two years across more than eighty candidates.

No indemnity for AI screening decisions

COST TO CLOSE: ONE LEGAL REDLINE PASS

Left open, there is nothing to recover from the vendor for defence or for damages.

01

Audit the vendor before you sign

These seven belong in every RFP for an AI hiring tool, and adding them takes about twenty minutes. If a vendor will not answer 1, 2 or 4, treat that as a disqualifying response rather than the opening of a negotiation.

- [ ] 01 The adverse impact analysis for the last twelve months, showing pass rates by race, sex, age over and under 40, and disclosed disability, measured against real customer applicant pools.
- [ ] 02 The retest cadence after any model update that could affect bias, with the results of the last three updates attached.
- [ ] 03 Every configuration parameter available to your own deployment team that can move a candidate's score or rank.
- [ ] 04 A description of the training data, including what share of it came from existing customers and the demographic composition of the historical hiring decisions inside it.
- [ ] 05 The process for notifying you of material model updates, with "material" defined in writing by them.
- [ ] 06 Any bias concern raised by a customer in the past twenty-four months, and how each one was resolved.
- [ ] 07 A contractual commitment to adverse impact testing, with the results delivered to you every year.

ALREADY SIGNED Ask them now, in writing. Whether they answer or stay quiet, that is what the record will show later.

02

Close the indemnification gap

The standard liability cap and the standard indemnity carve-out together mean the vendor owes you nothing if its tool produced a discriminatory result. Procurement usually misses this because it is reading the contract for uptime and data residency rather than for civil rights exposure.

	CLAUSE	WHAT YOU NEED INSTEAD	PUSH-BACK
☐	**Limitation of liability**	Uncapped for discrimination claims arising from the vendor's AI decisions	HIGH
☐	**Indemnification scope**	Covers claims arising from the vendor's screening decisions	VERY HIGH
☐	**Model update notice**	Thirty days written notice before any material update	MEDIUM
☐	**Bias testing**	An annual adverse impact analysis, delivered to you	LOW–MED
☐	**Configuration audit trail**	Your right to retrieve the full configuration history on request	LOW
☐	**Training data disclosure**	Annual disclosure of the demographic composition of the training data	HIGH

ABOVE ROUGHLY $500K A YEAR

Even the rows marked very high are negotiable at that value, so ask for them plainly and hold the position.

BELOW THAT, TRADE FOR LEVERAGE

Ask for mutual termination if an adverse impact threshold is breached, a price reduction tied to bias test results, or documentation of the vendor's own bias insurance. If your contract is signed with none of this in it, log it with legal as a known risk and negotiate it into the renewal.

03

Run the four-fifths rule on the data you already hold

The Uniform Guidelines on Employee Selection Procedures (EEOC, 1978) require any selection procedure to be validated against adverse impact, and an AI screen is a selection procedure like any other.

THRESHOLD

0.80

A group's selection rate divided by the highest group's rate. Anything below this figure is a breach you have to review.

- [ ] 01 Pull twelve months of applicants and screen pass-throughs, split by race, sex, age over and under 40, and disclosed disability.
- [ ] 02 Work out the selection rate for each group by dividing the number who passed the screen by the applicants in that group.
- [ ] 03 Identify the group with the highest selection rate and use it as your reference.
- [ ] 04 Divide every other group's rate by that reference rate and flag any ratio that comes out under 0.80.
- [ ] 05 If you find a breach, document it the same day, bring in legal, and establish whether it was there in prior periods as well.

EVERY QUARTER

Run it on a schedule and log the result each time, because that log is what stands as your documented monitoring record.

WHAT A BREACH LOOKS LIKE

The 0.67 and 0.50 ratios in this example are the pattern Mobley alleges. If numbers of this shape are sitting in a vendor dashboard and nobody acted on them, that inaction is the exposure.

GROUP	APPLICANTS	RATE	RATIO
White	1,200	30.0%	1.00 ref
Asian	300	29.0%	0.97
Hispanic	240	25.8%	0.86
Black	180	20.0%	0.67 fail
Over 40	400	24.0%	0.80 review
Disability disclosed	60	15.0%	0.50 fail

04

Put three checkpoints into the quarterly review

Monitoring an AI hiring tool is a data analysis task rather than a model performance problem, which means one person can carry it with three numbers to check. Budget under two hours if your applicant tracking system can export its reporting.

	CHECKPOINT	TRIGGER	WHAT YOU DO ABOUT IT
☐	**Screening stage**	Any four-fifths ratio under 0.80	Suspend the AI screen, bring in legal, and run the root cause before it resumes
☐	**Configuration drift**	Any undocumented change to scoring parameters	Document the change and require the sign-off retrospectively
☐	**Vendor model updates**	An update shipped with no documented retest	Request the updated bias analysis before the next screening cycle opens

IF YOU CANNOT

If your system cannot show rejection data by group at all, that absence is itself the finding, because you are running a screen you have no way to audit.

THE WHOLE CORRECTION

Every company in this case made defensible operational decisions with the information it had at the time. What none of them did was ask who was responsible for what the AI decided. Putting that question into the RFP, the contract review and the quarterly review is the whole of the fix.

If you operate in Europe, Annex III of the EU AI Act already classifies recruitment AI as high risk and requires conformity assessment, bias monitoring and human oversight before deployment, whichever way this case resolves in the US courts.

Sebastián Ocampo

SEBASTIÁN OCAMPO
AI SYSTEMS IN OPERATION
Y8Y.AI

REVIEWED BY

DATE COMPLETED