MβB

WHITE PAPER

The Signal Problem

Talent technology and AI in hiring: what has changed, what has not, and what to do about it.

WHERE BETTER BEGINS

September 2026

Position at 28 September 2026. Regulation and vendor claims move quickly. Check before you rely on anything here.

THE ARGUMENT IN ONE PAGE

Executive summary

Every hiring process runs on signals. A CV, an application, a covering note, a reply to an outreach message, an interview answer, a reference. Employers read those signals and decide whom to trust.

Generative AI has made signals cheap to produce and hard to verify. A candidate can generate a tailored application in seconds. An employer can screen, rank and message thousands of people in the time it once took to read fifty CVs. Both sides of the table now have industrial tools, and most hiring processes were designed for neither.

This paper makes five points.

1. **Adoption is running ahead of evidence.** Most talent leaders plan to use AI in 2026 and about half plan to add autonomous agents. Independent evidence that any of it improves quality of hire is thin, and much of what is published comes from the people selling the tools.
2. **Cheap signals are losing value.** Candidate fraud, AI-written applications and coached interview answers all weaken the things hiring has always relied on. Candidates know it, and only about a quarter trust AI to judge them fairly.
3. **Automation makes decisions whether or not you call it that.** The UK regulator found that many employers are probably relying on solely automated decisions in recruitment, and that several did not realise it.
4. **A human in the loop is only a safeguard if the human can change the outcome.** A reviewer who cannot override, or who defers to the machine, is a formality. Research shows people tend to mirror the bias of the recommendation in front of them.
5. **Regulation has shifted, not gone away.** The EU has moved its high-risk deadline for recruitment AI to December 2027. The UK regime is already live and the regulator has said recruitment is a priority.

The position we take

AI earns its place in hiring where it removes delay without judging anyone. It has to prove itself wherever it ranks or rejects people. And the organisations that benefit most will not be those with the most tools. They will be the ones that measured before they automated, kept judgement where it belongs, and built the capability to run their own process without a vendor holding their hand.

How to read this paper

We separate three kinds of statement. Evidence comes from a regulator, a court record or peer-reviewed research. Reporting comes from analysts, law firms and vendors, and we say which. MβB view is our professional judgement, informed by delivery experience. It is a position to test, not a finding.

1. Where things stand

Adoption has moved from experiment to default. The best available indicators:

84%
of talent leaders plan to use AI in recruiting in 2026
Korn Ferry TA Trends 2026 (reported)

52%
plan to add autonomous AI agents to their teams
Korn Ferry TA Trends 2026 (reported)

70%+
of organisations expect to increase AI and automation in recruitment over five years
ICO, March 2026

From assistive to agentic

The change that matters in 2026 is the move from assistive tools to agentic ones. An assistive tool suggests something when asked. An agentic tool plans and carries out a run of tasks towards a goal, such as sourcing candidates, sending outreach, screening replies and booking interviews, without a person triggering each step.

Deloitte's 2026 Human Capital Trends work frames the choice for talent acquisition as a move from ad hoc AI adoption to intentional AI design. We agree with the framing. Most of the risk in this market comes from adoption that was never designed.

The tools are arriving inside the tools you already own

Large HR platforms have been buying specialists in conversational recruiting and applicant management, with Workday and Paradox and SAP and SmartRecruiters the most cited examples. The practical effect is that AI features increasingly turn up inside systems employers already run. That matters because the ICO found that employers often did not recognise what their tools were doing (see section 4).

The honest gap

Adoption figures tell you how many people have bought something. They say nothing about whether it worked. Published impact claims, including faster time to hire and better retention, are overwhelmingly vendor-reported and rarely tied to a baseline set before the tool went live. We are not saying the claims are false. We are saying that most of them cannot be checked, and a business that cannot check a claim should not build a plan on it.

2. The signal problem

Hiring has always been a contest between signals and the truth. A polished CV was a modest signal of effort and competence because polish took time. A fluent interview answer was a signal of thinking because fluency was hard to fake in real time. Generative AI has cut the cost of producing both to almost nothing.

6%
of 3,000 candidates admitted taking part in interview fraud
Gartner survey, 2Q25 (self-reported)

1 in 4
candidate profiles predicted to be fake worldwide by 2028
Gartner prediction

26%
of candidates trust AI to evaluate them fairly
Gartner survey, 2025

Three things are worth noticing.

- **The 2028 figure is a prediction, not a measurement.** The 6% is a measurement, but it is self-reported, and people who cheat rarely volunteer it. Read it as a floor. Other fraud figures circulating in vendor marketing vary wildly, often because they come from a single firm's applicant pool. We have left them out.
- **Candidates are changing behaviour in response to being assessed.** Research cited by Gartner suggests that when people believe an AI is judging them, they distort how they present their skills to suit what they think the AI rewards. Measurement changes what is being measured.
- **Trust has a commercial cost.** Gartner also reported that only 51% of candidates accepted the offer in their most recent application, down from 74% two years earlier. We cannot say how much of that is caused by AI, and neither can Gartner. But a process candidates distrust is a process that loses people at the final step.

What breaks and what holds

MβB view. Signals that are cheap to generate and easy to submit at a distance are the ones losing value fastest: keyword-matched CVs, covering letters, unsupervised take-home questions, and unstructured video interviews. Signals that hold their value are those where the candidate has to demonstrate something, live or under consistent conditions, against criteria set in advance: work samples, structured exercises, and structured interviews scored against a rubric. None of that is new. What is new is that the case for doing it properly has stopped being optional.

Application volume compounds the problem. One applicant tracking vendor reports that auto-apply tools pushed application volumes up 93% in a year. Treat the figure as a vendor's, but the direction is consistent with what recruiters describe: more applications, less information per application. A screening process built for a manageable pipeline now faces a flood of near-identical, well-written submissions. Faster sorting does not fix this. It sorts the flood more quickly.

3. The other side of the table

Employers are automating too, and here the evidence is more uncomfortable.

What the research shows

In 2024, University of Washington researchers tested three large language models on how they ranked CVs. They took 554 real CVs and 571 job descriptions, swapped in first names associated with different races and genders, and changed nothing else. The results, from more than three million comparisons:

85%
of cases favoured white-associated names
Wilson and Caliskan, AIES 2024

11%
of cases favoured female-associated names
Wilson and Caliskan, AIES 2024

3m+
CV and job description comparisons tested
University of Washington

Black male-associated names were disadvantaged in up to 100% of cases. The caveat matters: the models tested were open-source retrieval models, not specific commercial products. The study does not prove any named vendor's tool behaves this way. What it shows is that the foundations many tools are built on can carry strong bias, which is exactly why you need the vendor's own test results rather than its reassurance.

A follow-up University of Washington experiment with 528 participants found that people choosing candidates without AI suggestions showed little bias, but when given biased recommendations they mirrored them. This is the finding that should worry anyone relying on "a human reviews the shortlist" as their safeguard. Review does not neutralise a biased recommendation. It can launder it.

The courts are working out who is responsible

In the United States, Mobley v Workday is the case employers and vendors are watching. The plaintiff alleges that the AI screening tools used across Workday's customers discriminated on grounds of age, race and disability. Workday denies the allegations and the merits are undecided. But the procedural rulings so far are significant. The court accepted in 2024 that a vendor performing screening functions could be treated as an employer for the purposes of federal discrimination law. It allowed an age-discrimination collective action to proceed nationwide in May 2025. And in June 2026 it refused to dismiss claims under California's Fair Employment and Housing Act.

MβB view. US law does not govern UK hiring, and the case may end in many ways. But the direction is clear enough for buyers. You cannot assume the vendor carries the risk of a tool you chose, configured and relied on. Contracts, testing and evidence trails need to reflect that.

4. The regulatory position

A summary of public sources. It is not legal advice, so take advice on your own position.

United Kingdom: live and focused on recruitment

The Data (Use and Access) Act 2025 rewrote the UK GDPR rules on solely automated decisions. The old near-blanket restriction has been replaced with a safeguard-led regime: automated decisions with significant effects are permitted, provided the right protections are in place, with the tightest limits reserved mainly for decisions relying on special category data.

On 31 March 2026 the Information Commissioner's Office published Recruitment rewired, a report based on voluntary engagement with more than 30 employers between March 2025 and January 2026, alongside draft guidance on automated decision-making. The regulator is clear that it supports responsible use. Its findings are still pointed:

- Many employers are **likely relying on solely automated decisions** in recruitment, meaning systems making decisions with significant effects on candidates without meaningful human involvement.
- Employers **need to be far clearer with candidates** about when and how automated decision-making is used.
- Where human involvement exists, it must be **applied consistently to all candidates at a stage**, not selectively.
- Employers should **expand monitoring for fairness and bias**.
- Most employers said their tools only supported decisions. The evidence suggested many were making substantive ones. Human involvement must be genuine: the reviewer needs the authority, discretion and competence to change the outcome before it takes effect.

The ICO wrote to 16 organisations it considered likely to be using automated decisions about candidates, and reports that they have committed to act on its recommendations. The consultation on the draft guidance closed on 29 May 2026 and the ICO indicated that final guidance would follow in summer 2026. Check the current status of that guidance before you rely on this summary.

European Union: delayed, not dropped

AI systems used for recruitment and selection are classed as high-risk under the EU AI Act. The original deadline for those obligations was 2 August 2026. The Digital Omnibus on AI, in force since 27 July 2026, moved it. The UK is outside the EU, but the Act applies where an AI system's output is used in the EU, so UK employers hiring into EU markets, and UK vendors selling into them, still need a view. Obligations for high-risk AI embedded in regulated products follow on 2 August 2028. The delay is sixteen months of breathing room, not a repeal: requirements for risk management, technical documentation and human oversight remain, and the sensible use of the time is to build those foundations.

Date	What happens
2 August 2026	General application of the Act begins, including the Article 50 transparency duties on AI interaction and AI-generated content.
2 December 2027	Obligations for stand-alone high-risk systems, including recruitment and selection tools, now apply. Previously 2 August 2026.

5. Where AI earns its place

MβB view. The table below is our professional judgement, drawn from the evidence above and from delivering hiring at scale. Use it to test your own process, not as a settled ruling. Where AI improves speed without judging a person, the case is strong. Where it ranks, scores or rejects, the burden of proof rises sharply.

Stage	Where AI helps	Where it creates risk	What people must own
Role definition and intake	Drafting role profiles, adverts and intake notes from existing material. Flagging inconsistent requirements.	Copies the flaws of its source: inflated requirements, biased wording, generic adverts. A faster route to the wrong role.	What the job really is, what it pays, and what good looks like.
Sourcing	Widening search, finding adjacent skills, keeping pipelines warm. Strong for repeatable, high-volume search.	Rewards profiles that resemble past hires. Inferring traits the candidate never supplied. Spam that damages your reputation.	Defining the search. Checking the shortlist against reality.
Outreach and engagement	Personalised first drafts, instant answers to routine questions, scheduling.	Volume without relevance. Undisclosed AI conversations. Candidates who cannot reach a person.	Honesty about the role. Any conversation that shapes a candidate's decision.
Screening and shortlisting	Consistent first-pass sorting against explicit, job-related criteria at volume.	The highest legal and bias exposure. Scores that cannot be explained. Probably a solely automated decision.	Setting and testing the criteria. Reviewing rejected as well as shortlisted candidates. Owning every rejection.
Assessment and interview	Structured question banks, transcription, note-taking, consistent scoring rubrics.	Behavioural and video analysis draws public distrust and special category risk. Take-home tasks are now answerable by AI. Candidates game what they think is measured.	Deciding what evidence counts. Running the conversation. The judgement.
Identity and verification	Document checks, liveness detection, anomaly flags in remote hiring.	False positives penalise honest candidates. Detection is an arms race. Checks at the end arrive after the decision has formed.	Risk-based rules. Handling flags fairly. Escalation.
Scheduling, admin and onboarding	The strongest case. Removes delay without judging anyone. Reminders, documentation and early check-ins for new starters.	Badly configured automation can drop candidates silently. New starters can feel processed rather than welcomed.	Owning exceptions. The relationship and early-life support.
Decision and offer	Summaries and comparison packs for decision-makers.	Anchoring on the machine's summary. Automation bias.	The decision, and the reasons for it.

6. Six principles for integrating AI

MβB view. None of these are complicated. Very few organisations do all six.

1. Set the baseline before the tool arrives

Record cost per hire, time to fill, conversion at each stage, offer acceptance, candidate drop-out, hiring manager satisfaction and outcomes by demographic group where you lawfully can. If you did not measure before, you cannot claim after, and you cannot tell a vendor's promise from a result. Decide in advance what result would make you switch the tool off.

2. Automate the admin, protect the judgement

Scheduling, note-taking, data capture and first drafts have a strong case. Ranking and rejecting people carries a heavier burden. Start where the risk to candidates is lowest and the benefit is easiest to measure, and earn your way up.

3. Make human involvement able to change the outcome

Give reviewers the information, the authority and the time. Apply review to every candidate at that stage. Record overrides. If reviewers never disagree with the tool, that is a reason to ask whether review is real, not evidence that the tool is perfect.

4. Redesign assessment for a world where answers are free

If a chatbot can produce a strong answer to your screening question, the question no longer tells you what you think it does. Shift weight to structured, job-related exercises where candidates demonstrate something under consistent conditions, and verify identity before the evidence is weighed rather than after.

5. Buy the evidence, not the demo

Ask every vendor for their bias testing method and results, what is switched on by default, what data trains their models, who carries which risk in the contract, and how they support candidate rights. A vendor who cannot answer clearly has answered.

6. Build the capability that outlasts the tool

Tools change every eighteen months. Trained recruiters, hiring managers who can run a structured interview, and a named owner with authority to pause a tool do not. Avoid a position where the only people who understand your hiring process work for your software supplier.

7. What we do not yet know

Where the evidence runs out, we say so.

- **Whether AI improves quality of hire.** Speed and cost claims are common. Robust, independent evidence on who performs well and stays is not.
- **How the ICO's final guidance will read**, and how firmly the regulator will enforce. The signals so far are strong.
- **Whether agentic tools sustain performance** beyond repeatable, high-volume tasks, and what happens to candidate trust as they spread.
- **How courts and regulators will divide responsibility** between employers and vendors, and whether candidate trust recovers from its current low.

What to do this quarter

1. **List every tool in your hiring process that uses AI or automated scoring,** including features inside your existing ATS and job boards. Ask each vendor in writing what is switched on. Use the companion checklist.
2. **Pick one stage and set a baseline.** Use the decision record template to state the problem, the measures, the human role and the conditions under which you would stop.
3. **Test one vendor claim properly.** Use the vendor question set. Score the answers, not the demo.
4. **Decide who owns this.** One named person, with authority to pause a tool.

Companion resources

AI in Hiring Assurance Checklist: the questions to answer before you switch a tool on, and every quarter after.

AI Hiring Tool Decision Record and Vendor Question Set: an editable template for each use case, and a scored set of questions for comparing suppliers.

About MβB

MβB is a talent and business consultancy that tests accepted thinking, examines what is actually happening and improves the systems connecting people with performance. It is led by Mike Bratby-Bale, who has spent more than twenty years running recruitment in-house across nuclear, defence, infrastructure, utilities, financial services and logistics. That includes delivering 721 hires in five months on a major nuclear infrastructure programme and reducing agency spend by 68% in a single year.

Does your hiring pass the Beta Test? If you would like an independent view of how AI and automation are working in your process, speak to us through the contact page of the site you downloaded this from.

Sources and notes

Type: R = regulator or court record; RS = peer-reviewed or university research; A = analyst; L = law firm commentary; V = vendor or secondary report. Where we cite a figure reported second-hand, we have said so in the text.

Type	Source
R	Information Commissioner's Office. Recruitment rewired: an update on the ICO's work on the fair and responsible use of automation in recruitment. 31 March 2026. ico.org.uk/about-the-ico/what-we-do/recruitment-rewired/
R	ICO. Automated decisions can streamline the hiring process, with the right safeguards in place. March 2026. ico.org.uk (news and blogs, March 2026).
L	DLA Piper, Technology's Legal Edge. UK: ICO report on automated decision-making in recruitment. April 2026.
L	Bird & Bird. ICO launches consultation on draft guidance on automated decision-making and profiling. May 2026.
L	Covington, Inside Privacy. UK ICO consults on draft automated decision-making guidance and sets expectations for ADM in recruitment. April 2026.
V	The Leveraged Years (TLY). UK ICO's AI decision guidance targets hiring. July 2026. Consultation dates and DUAA summary.
L	Gibson Dunn. EU AI Act Omnibus agreement: postponed high-risk deadlines and other key changes. May 2026.
L	Kinstellar. The AI Act after the Digital Omnibus. 2026.
L	Jones Walker. Yes, August 2 still matters. July 2026.
V	ActuIA. HR tools and artificial intelligence: Europe delays high-risk obligations to December 2027. May 2026.
V	Mission Connected, The Nonprofit Recruiter. The EU AI Act deadlines moved. What that actually means for hiring. 2026. Publication and entry-into-force dates.
A	Gartner. Survey shows just 26% of job applicants trust AI will fairly evaluate them. Press release, 31 July 2025.
A	HR Dive. By 2028, 1 in 4 candidate profiles will be fake, Gartner predicts. August 2025.
RS	Wilson, K. and Caliskan, A. Gender, race, and intersectional bias in resume screening via language model retrieval. AAAI/ACM Conference on AI, Ethics, and Society (AIES), 2024.
RS	University of Washington News. AI tools show biases in ranking job applicants' names according to perceived race and gender. 31 October 2024. Also the follow-up study with 528 participants, listed on the UW News page for Kyra Wilson.
A	SHRM. The Workday AI lawsuit is a wake-up call for HR. July 2026.
R	Civil Rights Litigation Clearinghouse. Mobley v. Workday, Inc., No. 3:23-cv-00770 (N.D. Cal.).
A	Deloitte. 2026 talent acquisition technology trends: the new imperative. May 2026.
V	Recruiterflow. The state of AI in recruiting 2026. August 2026. Reporting Korn Ferry, 12th Annual Talent Acquisition Trends.
V	Noon. AI in talent acquisition: what it actually changes in 2026. August 2026. Reporting Gem 2026 application volume data.
V	Softwarefinder. AI recruiting trends in 2026. Reporting platform acquisitions (Workday and Paradox; SAP and SmartRecruiters).