Kyle & Co

Kyle & Co · November 2026

The Raw Materials: Early-Career Hiring in the AI Era

The signals leading employers rely on when the resume can't be trusted

By Kyle Lagunas, Principal Analyst



Sponsored by CoderPad

Contents

This research was sponsored by CoderPad, whose team contributed to the survey design and reviewed an early read of the findings. The analysis, the conclusions, and this report are Kyle & Co's alone.

Introduction 3

Executive Summary 8

Section 1: The Signal Problem 10

Section 2: Where the Signal Comes Back 21

Section 3: How Leaders Differ from Laggards 27

Close the gap 37

Our Research Methodology 39

About Kyle & Co · About CoderPad 40

Introduction

Early Careers' Big Problem: Betting on the Future with Evidence of the Past

Hiring has always been a bet on the future made with evidence about the past. For experienced candidates, that past is full of referenceable information you can assess, test, and validate. There's a work history to verify, a former manager to call, shipped work to point at. The evidence may not always be perfect, but it's there.

In comparison, early career candidates don't have as much to go on. A limited (relevant) employment history. No former manager who worked with them beyond an internship. No body of work beyond school work or ad hoc pet projects. Employers are betting on what someone might become based on a thin record of what they've already done.

AI made that thin record even thinner. The documents candidates submit–from resumes and cover letters to written responses and take-home work–can now be generated, polished, and tailored to a job posting in seconds.

The problem isn't that candidates are cheating; in fact, most are doing exactly what they've been told to do everywhere else: use the tools. The problem is that employers can no longer tell what a document says about the person who submitted it.

Every kind of hiring is losing signal. Early-career hiring is the sharp end, because there's nothing underneath the documents to fall back on.

Raw materials: What early-career hires bring doesn't show up on paper

An early-career hire doesn't arrive with the skills of the role. That isn't a flaw. It's the definition. What they bring are the raw materials: how they think through a problem they haven't seen before, how fast they learn, how they handle being wrong, and increasingly, how well they work alongside AI and know when not to trust it.

Those raw materials are real, and they predict a great deal about who someone becomes. But they don't show up on paper. They show up in the work. Recruiters are starting to reach for new language to describe what they're looking for.



> "How do we hire for potential and slope? People who can adapt in a new environment and grow with it, and have agency to move and change quickly. What are the kinds of questions we ask to assess for that?"
>
> — Vivian Chen, Principal Recruiter, Lyft

The question this report tries to answer is whether employers have found a way to see them.

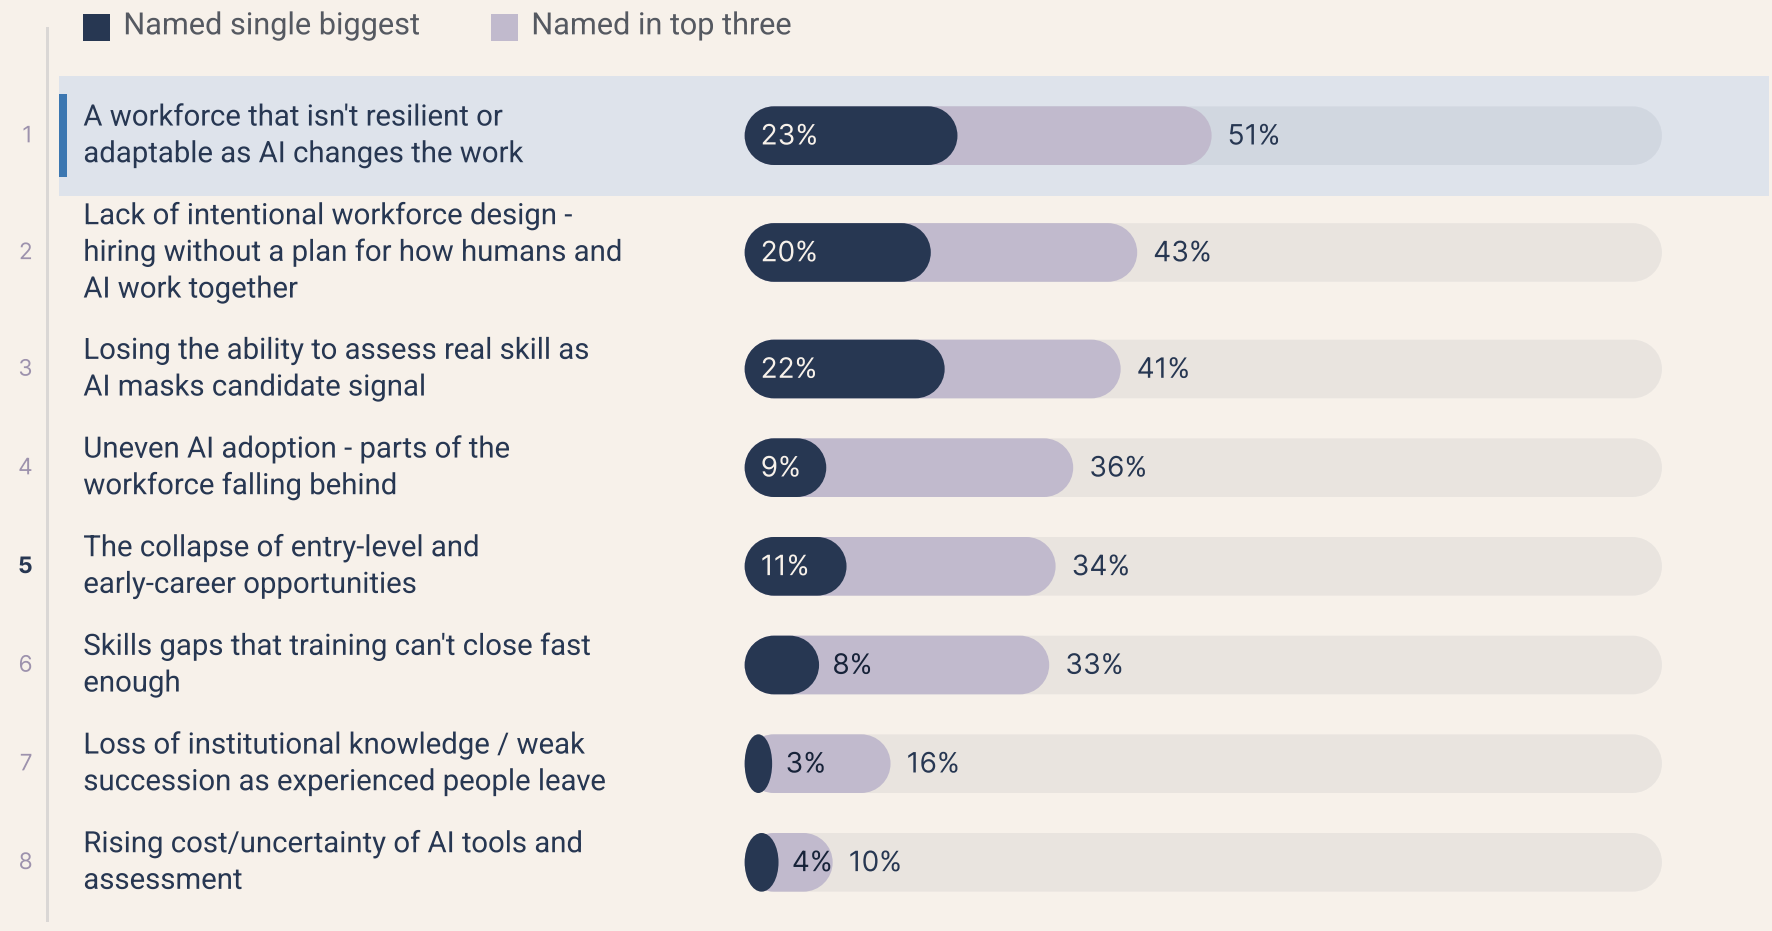
Junior hiring isn't going away. How far the bar should reach is still unsettled

It would be easy to read the headlines about AI and entry-level work and conclude that early-career hiring is winding down.

The employers we surveyed don't see it that way. Asked to name the biggest risks to their talent and workforce, they ranked the collapse of entry-level opportunity fifth of eight. It's a real concern, but it sits well behind worries about resilience, workforce design, and losing the ability to assess real skill. Employers still plan to hire early-career talent. What they're less sure of is how to tell who's ready.

Figure 1

Biggest risks to talent and workforce



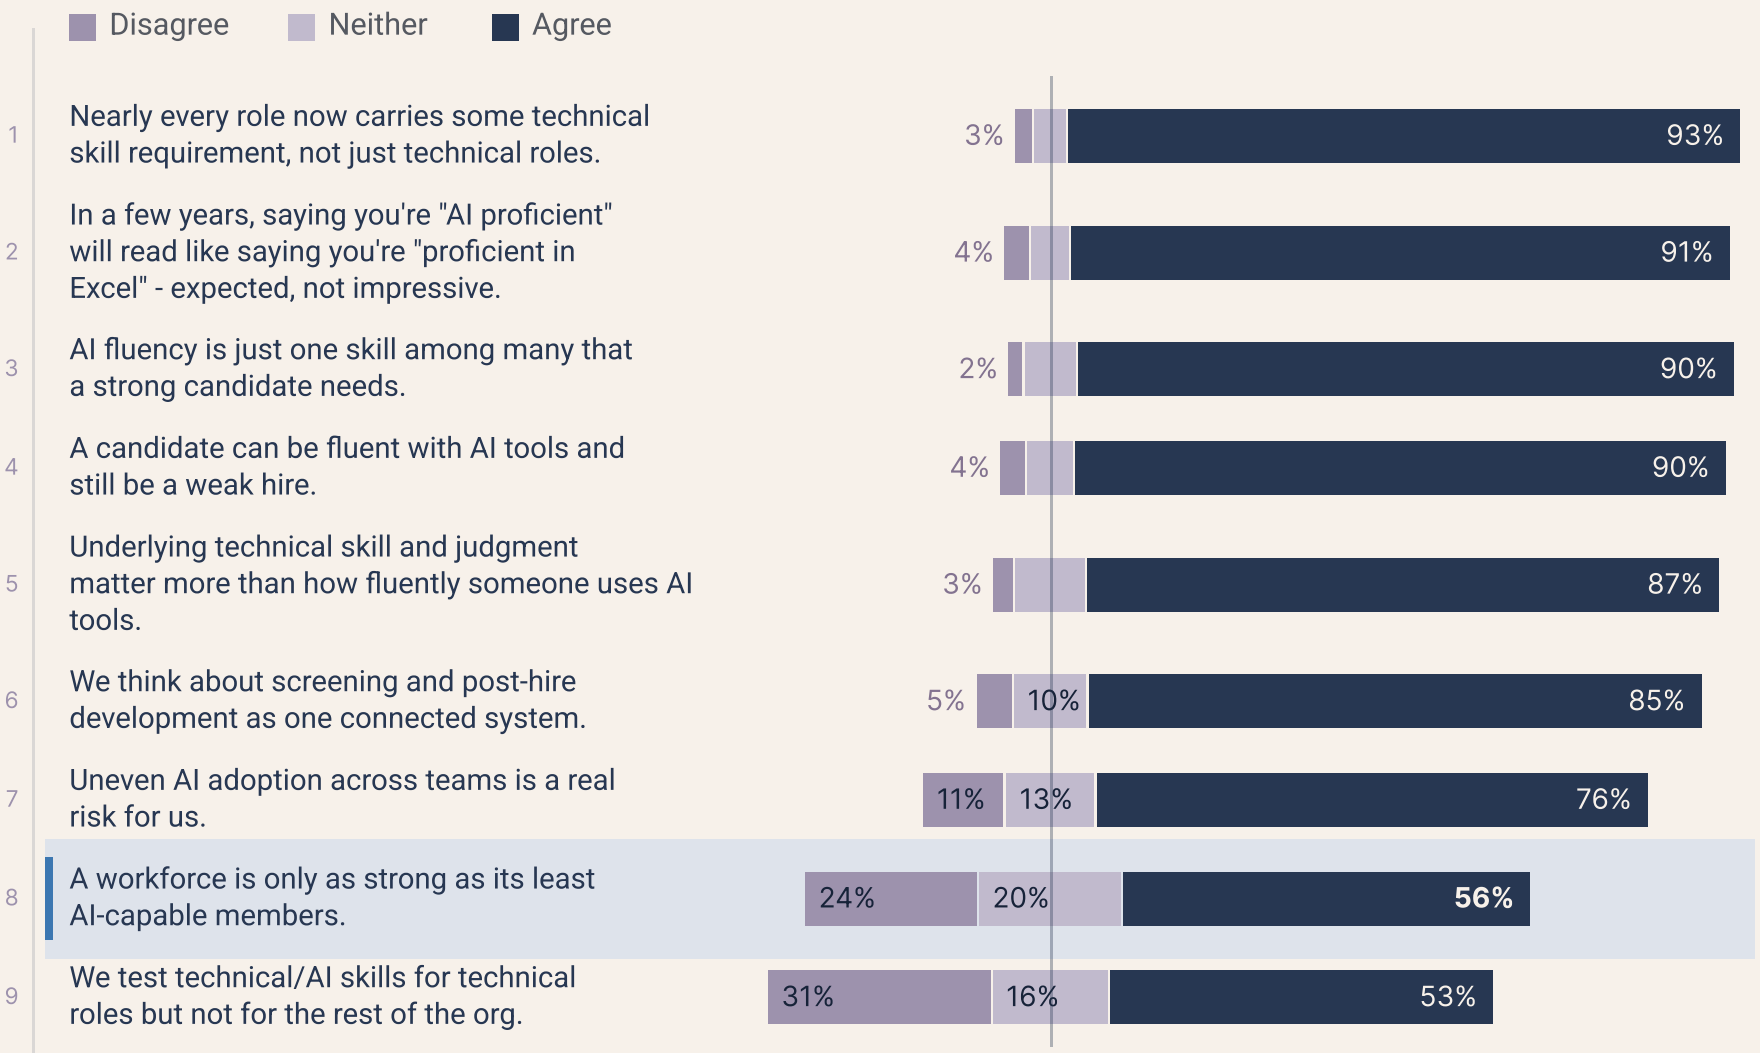
Share of all employers naming each risk in their top three, and as their single biggest risk.

That uncertainty shows up in one statement in particular. We asked whether "a workforce is only as good as its worst AI user." Of every attitude statement in the survey, it was the only one where the market

genuinely split. More than half agree (56%). Nearly a quarter reject it outright (24%). Everywhere else, employers largely agree with one another: AI proficiency will soon be table stakes, underlying skill and judgment matter more than tool fluency, and nearly every role now carries some technical requirement.

Figure 2

Attitudes toward AI, skill, and assessment



Ordered by share who agree. Neither is centered on the midline.

The split tells us where the unresolved question actually lives. It isn't whether AI capability matters. It's how far down the organization the bar should reach, and therefore how far assessment should reach. That is a hiring question before it is anything else.

About the Research

Most research on hiring in the AI era asks what employers believe. This study asks what they do, then checks it against how their hires turn out.

We built it to test the market's assumptions, not to confirm them. We anchored every finding in self-reported outcomes: quality of hire, first-year retention, and speed to productivity. And we're validating what the numbers show with the practitioners who do this work every day. That's the Kyle & Co standard, and it matters more here than usual. The question at the center of this report is one where everyone already has an opinion, and very few have evidence.

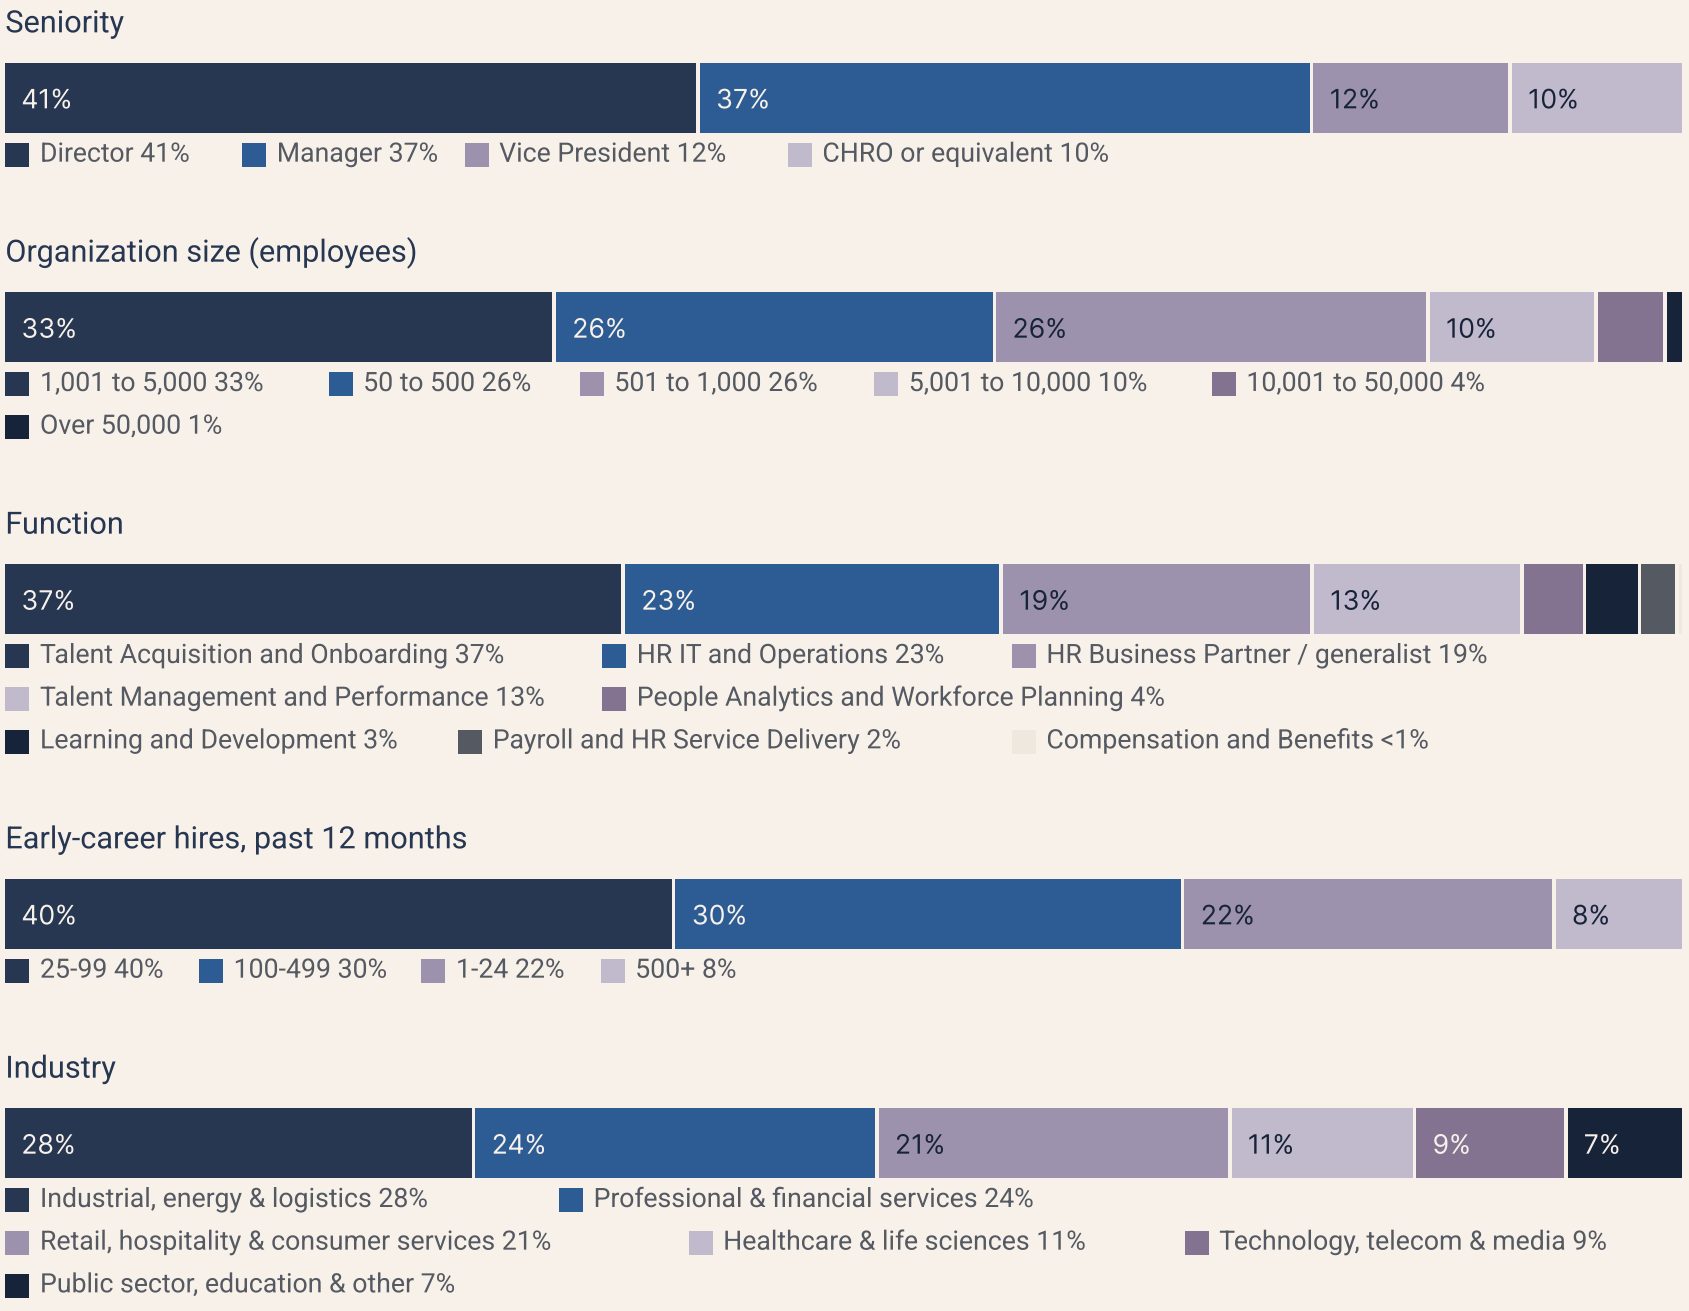
Who we heard from: Leaders accountable for early-career hiring

We surveyed 271 HR and talent acquisition leaders who are involved in or accountable for early-career hiring.

Respondents range from managers to CHROs, and more than three in five hold director-level roles or above. They work at organizations from 50 employees to more than 50,000, across retail, manufacturing, professional services, healthcare, technology, financial services, logistics, hospitality, education, and more. Most are hiring early-career talent at meaningful volume: nearly eight in ten made 25 or more early-career hires in the past year.

Figure 3

Who we heard from



Share of sample. Each bar sums to 100%.

Alongside the survey, we interviewed talent leaders who run early-career programs, including customers of CoderPad. Kyle & Co conducts every interview, and transcripts are not shared with the sponsor. Full details are in the Methodology.

How we define leaders and laggards

Throughout the report, we compare two groups defined by outcomes, not by the practices we go on to study.

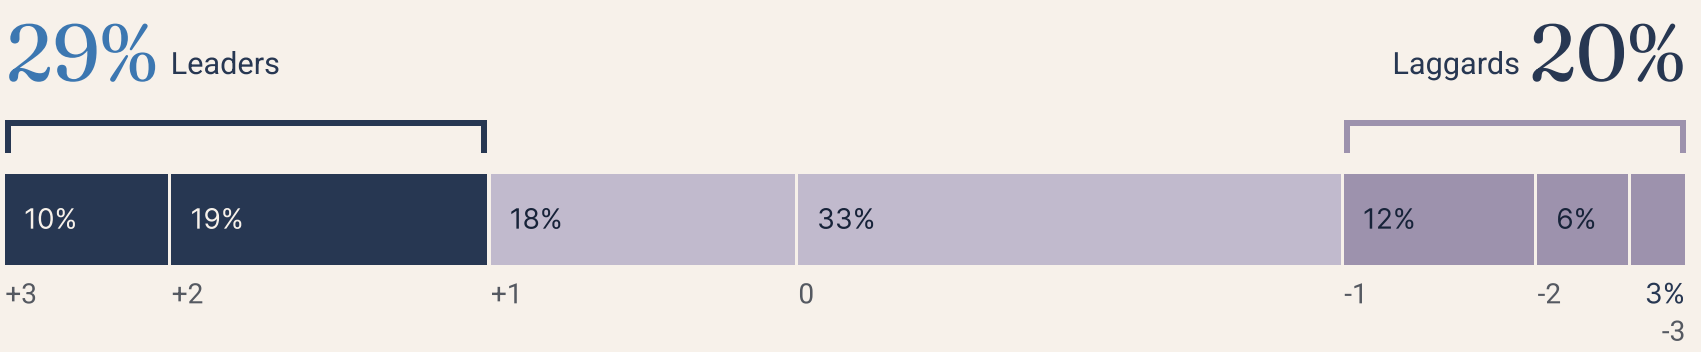
Leaders (29% of employers who track all three measures) rate themselves ahead of peers on early-career outcomes across quality of hire, first-year retention, and speed to productivity.

Laggards (20%) rate themselves behind.

These are self-assessments, and we treat them that way. Every gap we report is an association between what employers do and how they say they're performing. None of it proves one causes the other. What it does show is where the market's best-performing employers have made different choices than the rest, and those choices turn out to be specific, observable, and within reach.

Figure 4

How leaders and laggards are defined



Outcome score across the three early-career KPIs. Base: employers measuring all three (93% of sample). Middle group (+1, 0) muted.

Executive Summary

Employers have lost the ability to see whether early-career candidates can do the job. The ones winning at early-career hiring are the ones who stopped reading about candidates and started watching them work.

The argument moves from the fear employers hold, to the instruments they use, to what the best of them built instead. Six findings carry it.

The State of Early Career Hiring: Six Major Findings

Major Finding 1: The thing employers fear most changes nothing about how they hire

Half name a workforce that can't adapt as AI changes the work as their top risk. Yet employers who worry about it screen for adaptability at the same rate as those who don't, because a second risk, losing the ability to see real skill, stands between the belief and the practice.



Major Finding 2: Employers screen for potential using tools that measure history

More than half of what employers screen for describes potential, but the resume is still the first tool they reach for. One in three never watches an early-career candidate do the work before making an offer.



Major Finding 3: Early-career is where this becomes solvable

Quality of hire is what this market prioritizes most and spends on most. With no work history to fall back on, early-career is the hardest case, and a process that works without a track record works everywhere.



Major Finding 4: Almost no one is assessing the two things that are new

AI fluency and judgment rank last on what employers screen for and last on their confidence in assessing them, even though nearly nine in ten say judgment matters more than tool fluency.



Major Finding 5: The resume lost its relevance to ambiguity not to AI.

Employers who give candidates clear rules on AI use get strong signal from the resume more than twice as often as employers who set none.



Major Finding 6: Employers who watch candidates work are ahead

The share of employers ahead of peers on quality of hire climbs with each observed method they use, and doesn't move with methods completed away from observation.



A Clear Difference in Approach & Outcomes: Leaders versus Laggards

Behind those findings are five practices that separate leaders from laggards, which we examine in Section 3. Leaders don't see the problem more clearly than anyone else, and they aren't further along on AI. They worry about the same things in the same proportions. The difference is what they built, most of it unglamorous, and the gap between the top and the bottom of the market comes down to about one practice.

None of this requires a new philosophy of hiring. It requires employers to look at the work instead of the paperwork, and to tell candidates the rules before asking them to follow them.

Section 1:

The Signal Problem



Finding 1: The thing employers fear most changes nothing about how they hire

The risk employers name most often has almost no effect on how they hire. Asked to pick the biggest risks to their talent and workforce, half chose a workforce that isn't resilient or adaptable as AI changes the work. No other risk comes close in breadth.

> **Half of employers fear a workforce that can't adapt. The fear never reaches the screen.**

You'd expect that fear to show up in screening. It doesn't.

Employers worried about resilience screen for adaptability at the same rate as employers who aren't worried at all, just under three in five either way. Learning agility tells the same story: a near-identical share of both groups name it among the five things they screen for most. The fear is real. It just never reaches the hiring process.

Losing the ability to see real skill is why the fear stalls

A second risk explains why.

Two in five employers name losing the ability to assess real skill as AI masks candidate signal. Fewer employers pick it, but the ones who do hold it harder: more than half of them call it their single biggest risk, a higher conversion than any other risk in the study. (Conversion figures in Figure 1.)

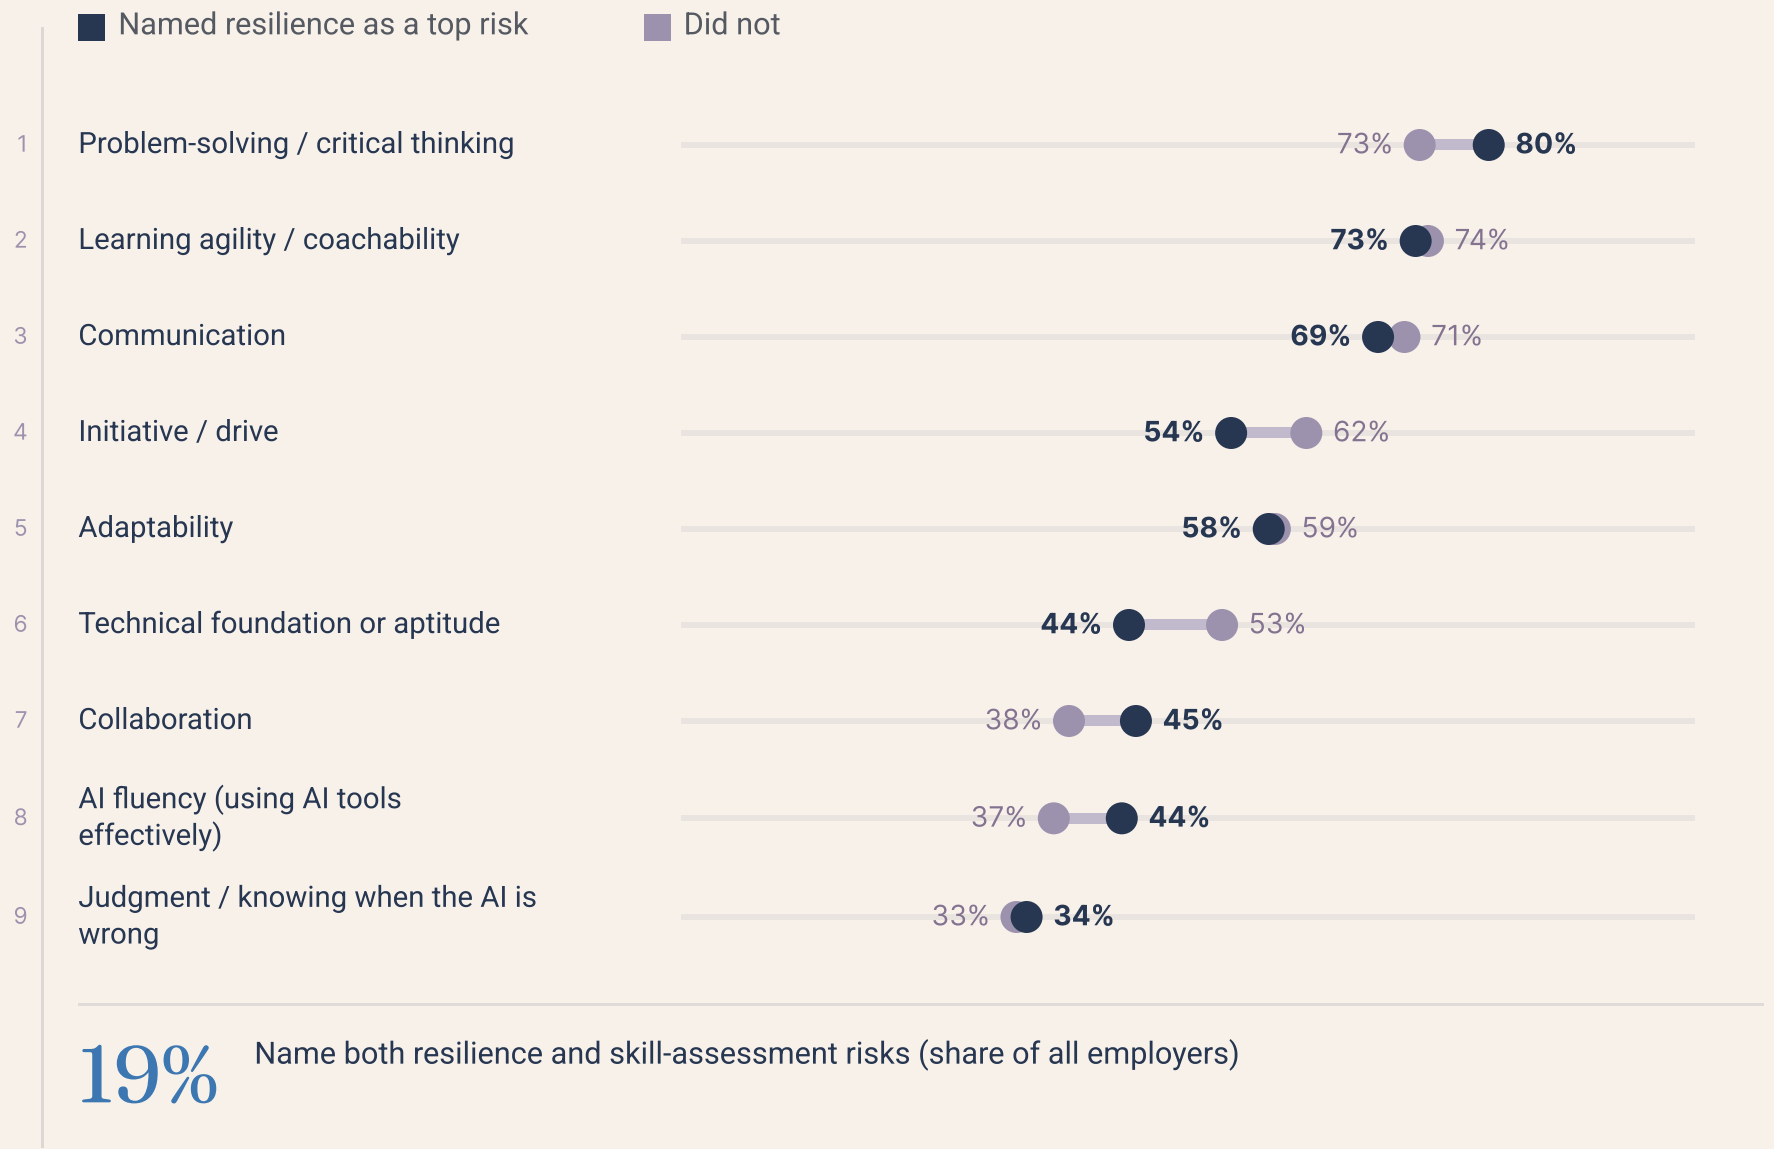
These look like two different worries, but they're actually one problem seen from two ends:

- [x] Resilience is what employers want to hire for
- [x] The inability to see real skill is why they can't.

An employer can believe adaptability matters more than anything and still have no reliable way to tell whether the candidate in front of them has it. So the belief stays a belief.

Figure 5

The fear doesn't change the screen



19% Name both resilience and skill-assessment risks (share of all employers)

Share naming each quality among the five they screen for most, by whether they named resilience as a top risk.

Only one in five employers see both the goal and the obstacle

Only one in five employers name both risks. Most see the goal or the obstacle, rarely both. That gap is where this report starts: the market knows what it needs from early-career talent, and it has lost the instruments to see it.



"People are still trying to figure it out. Right now it is a transition period where there's a lot of experimentation happening. We're starting to get some structure and guidance from leadership about what is and is not allowed. But if I zoom out, this is largely a period of figuring things out and experimentation."

— Vivian Chen, Principal Recruiter, Lyft

Finding 2: Employers screen for potential using tools that measure history

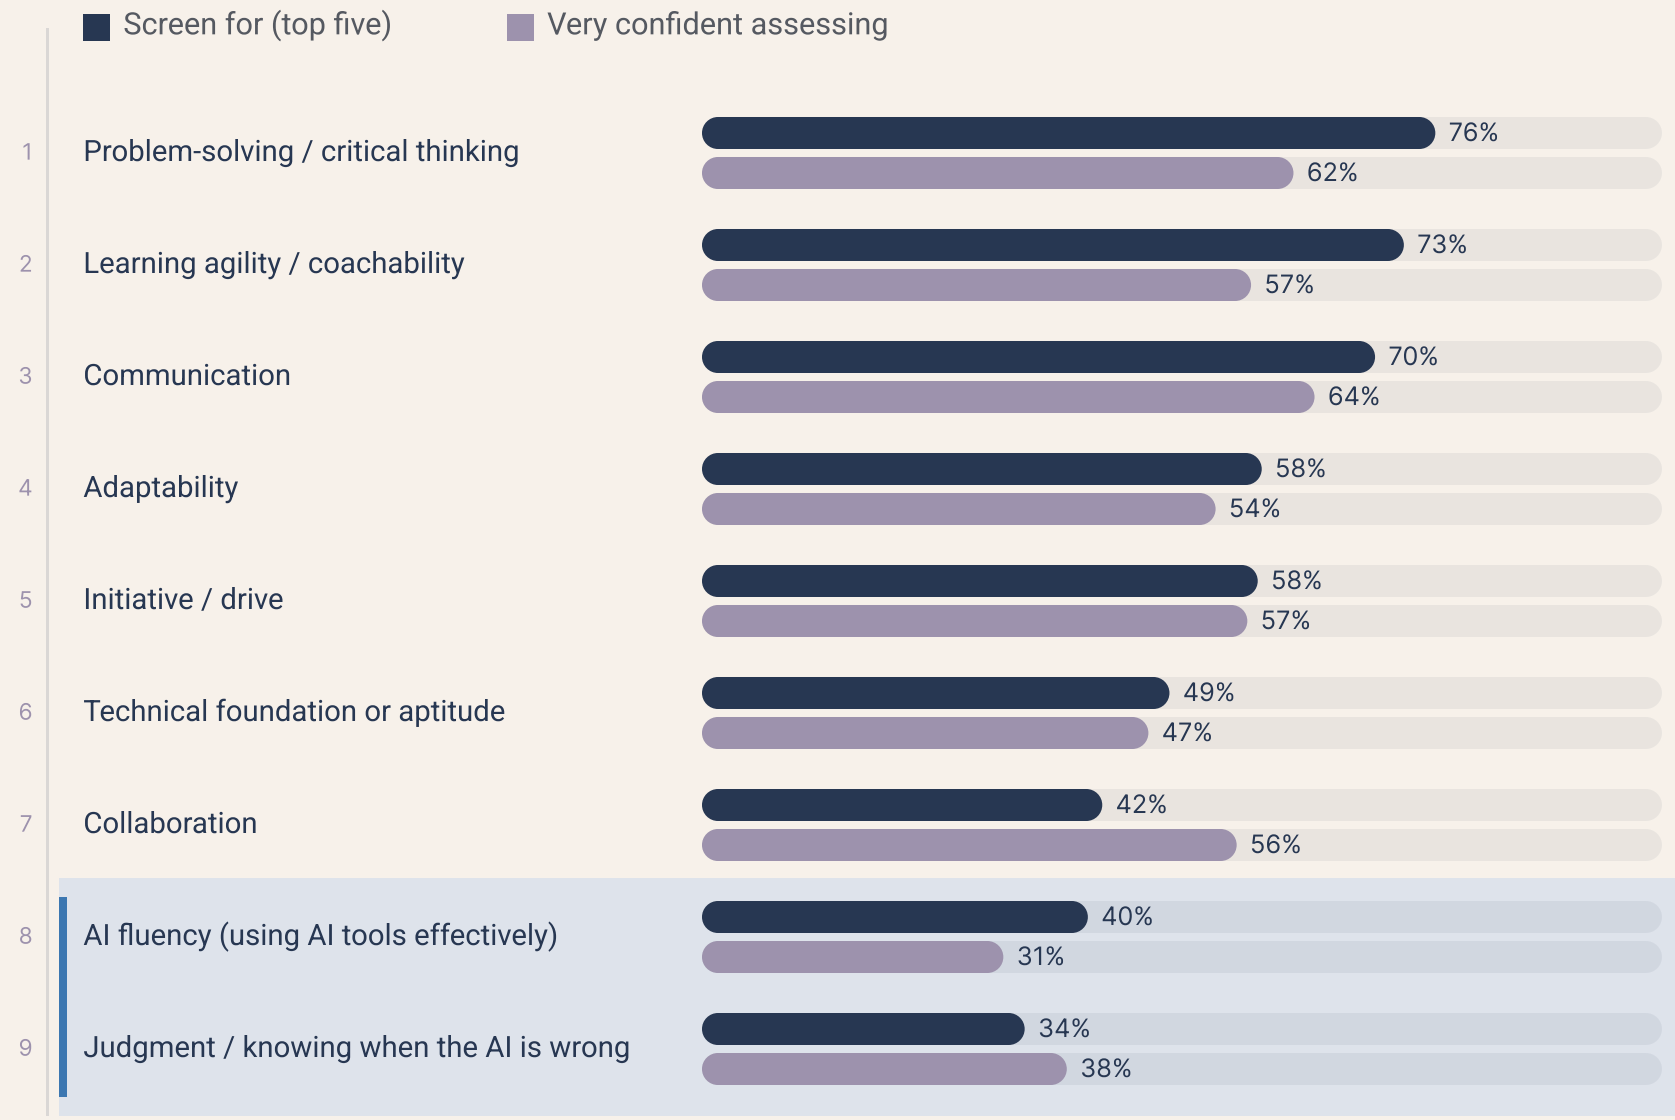
The mismatch at the heart of early-career hiring is simple to state. Employers are looking for potential, and the tools they use mostly measure history.

More than half of what employers screen for is potential, not accomplishment

Look at what employers say they screen for (Figure 6). The top three are problem-solving, learning agility, and communication, each named by seven in ten or more. More than half of everything on the list describes potential rather than accomplishment: how someone thinks, learns, and adapts, not what they've already done.

Figure 6

What employers screen for, and how confident they are assessing it



Share naming each quality among their top five, and share very confident assessing it.

Tell, talk, do: the "do" category is where potential shows, and it's thin

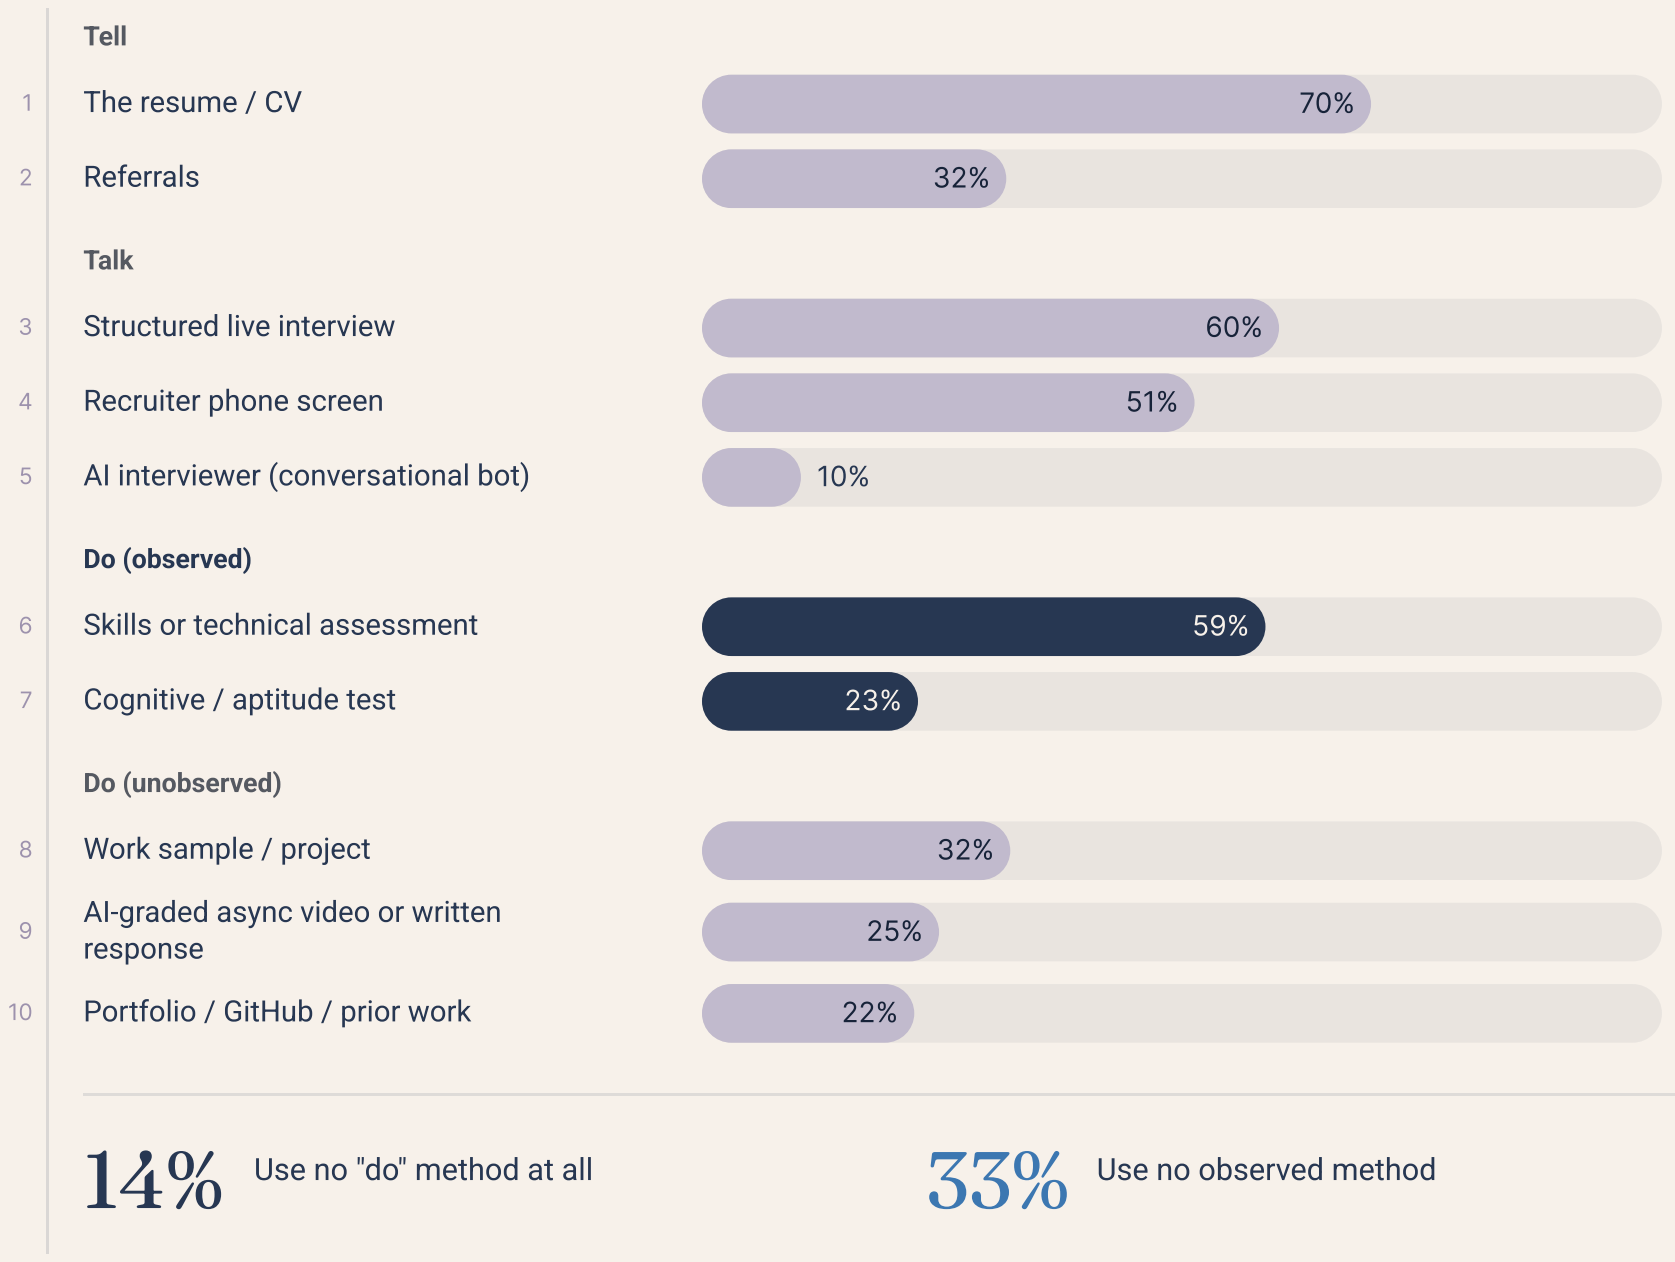
Now look at the instruments. The first one employers reach for is the resume, used by seven in ten, ten points ahead of anything else. A useful way to sort the rest is by what the candidate is asked to do:

- Tell: documents about the candidate, like the resume and referrals.
- Talk: conversations, like the recruiter screen and the structured interview.
- Do: the candidate produces or performs something, like a skills assessment, work sample, cognitive test, or portfolio.

The "do" category is where potential actually becomes visible, and it's thin. A general skills or technical assessment is the one exception, used by nearly three in five. Every other "do" method is used by fewer than a third of employers. Portfolio and prior work, the closest thing an early-career candidate has to a track record, is used by barely one in five.

Figure 7

Tell, talk, do: how employers assess early-career candidates



14% Use no "do" method at all

33% Use no observed method

Share of employers using each method with early-career candidates.

For a third of the market, nobody watches the candidate work before the offer

The result: one in seven employers hire early-career talent without ever asking the candidate to do anything. Widen the lens to whether anyone actually watches the work happen, and the share rises to one in three. For a third of the market, no one ever sees an early-career candidate do the work before making the offer.

That was a manageable gap when documents were hard to fake. It isn't anymore.



> "That's how most recruiters recruit. You look at the resume. What has this person done? Is this related to what we're doing? But how do you look for that highly adaptable, high-slope type of person?"
>
> — Vivian Chen, Principal Recruiter, Lyft

Finding 3: Early-career is where the resume signal problem becomes solvable

If the signal problem runs through all of hiring, why focus on early-career? Because it's where the stakes are clearest and the fix is most transferable.

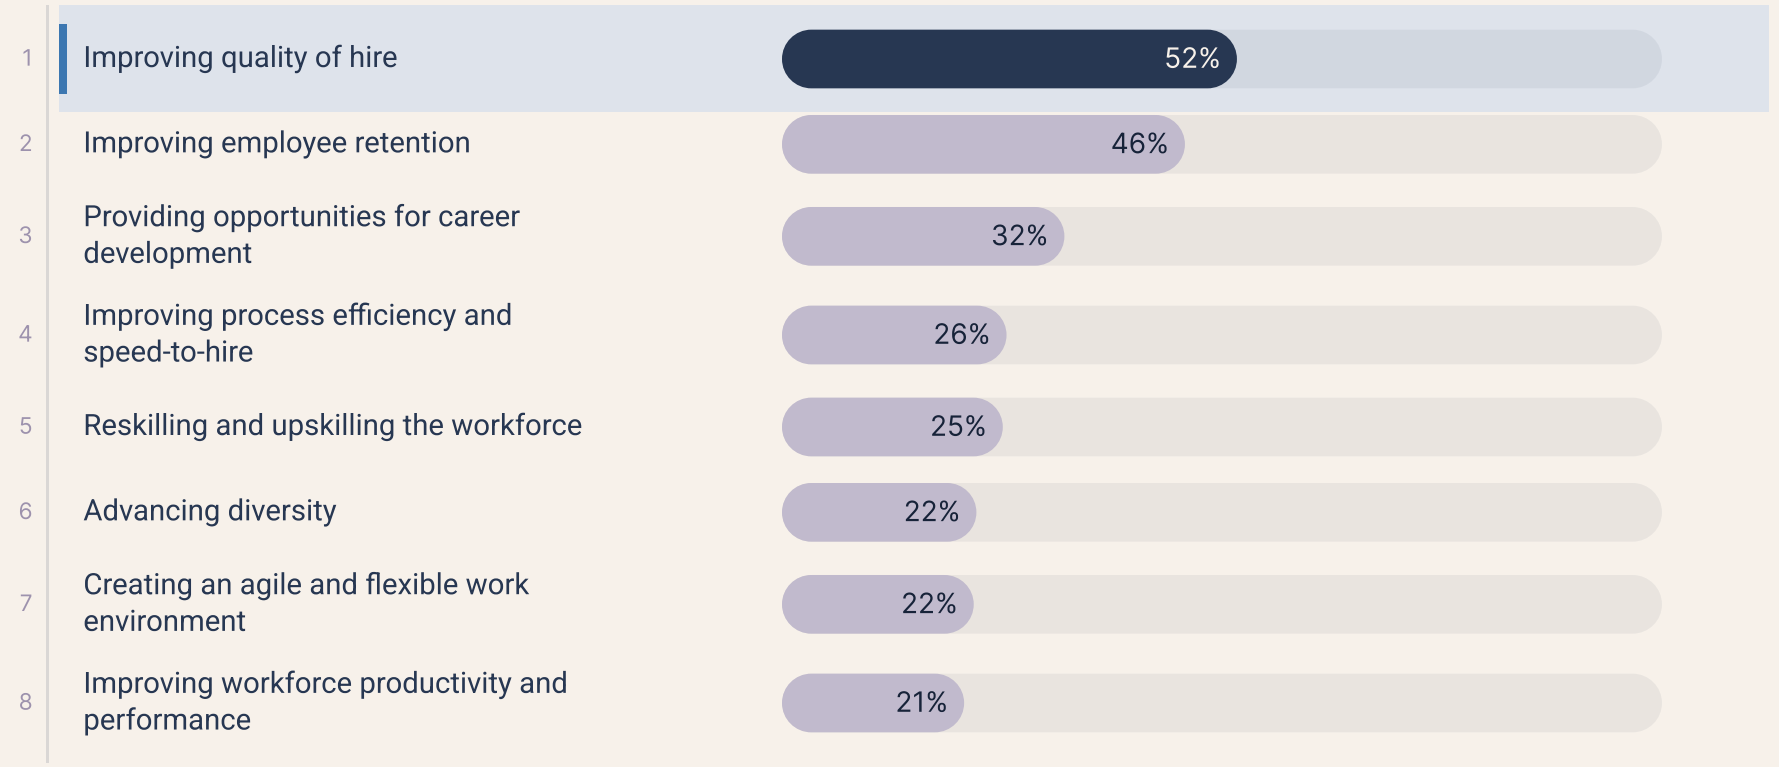
Quality of hire is what this market prioritizes and pays for

Start with what the market cares about. Improving quality of hire is the top workforce priority of the year, named by more than half of employers, ahead of retention and well ahead of speed. When employers justify spending on screening or AI hiring tools, quality of hire is the reason nearly three-quarters of them give.

Nothing else comes close. This is a market that has already decided what it's optimizing for.

Figure 8

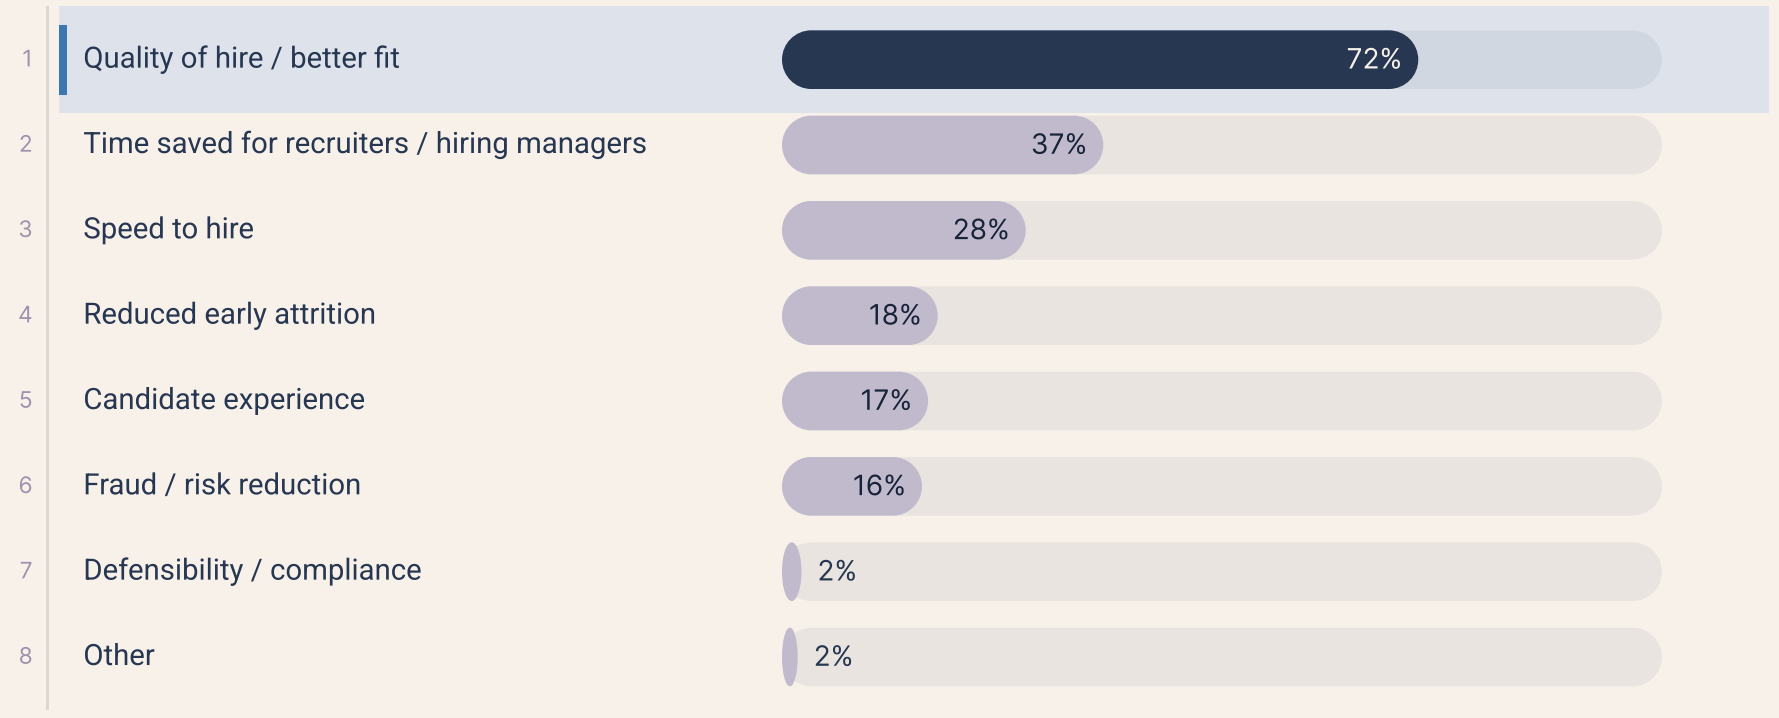
Top workforce priorities this year



Share naming each among their top workforce priorities this year.

Figure 9

What justifies spend on screening and AI hiring tools



Share naming each as a justification for spend on screening and AI hiring tools.

Without a track record, the process has to carry the whole decision

Early-career hiring is where that goal is hardest to reach.

With experienced candidates, a weak screening process gets bailed out by the track record: a decade of work history covers for a lot. With early-career candidates, there's nothing to bail it out. The process has to carry the whole decision.

That makes early-career the hardest case, and the most useful one. A process that can see a candidate's raw materials without a track record will work for anyone. A process that relies on the track record will quietly fail wherever the record is thin or, increasingly, generated.

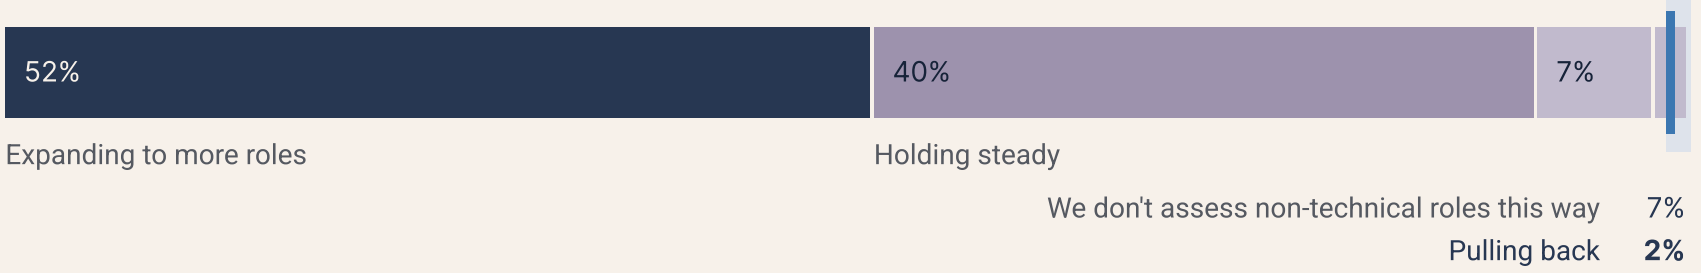
More than half are extending assessment. Just 2% are pulling back.

The appetite to fix it is already there. More than half of employers are extending skills-based assessment to more roles beyond technical ones. Just 2% are pulling back. The direction of travel is settled.

The open question is what that assessment looks like–and whether anyone is watching when it happens.

Figure 10

Direction of skills-based assessment



Share of employers, by the direction of skills-based assessment for non-technical roles.

Box is using early-career hiring as the testing ground for how it will assess everyone else.



"Our early-career roles were first, and now we're slowly rolling out to our full-time roles."

— Megan Yip, University Recruiting, Box

Finding 4: Almost no one is assessing the two things that are new

The two capabilities AI made newly important are the two employers screen for least and trust themselves least to judge.

AI fluency and judgment rank last on screening and last on confidence

Of the nine qualities we asked about, AI fluency and judgment (knowing when the AI is wrong) rank last and second-to-last on what employers screen for in early-career candidates. They also rank at the bottom on confidence: fewer employers feel very confident assessing either one than any of the other seven. (See Figure 6 above)

Belief isn't the problem: 93% say nearly every role now carries a technical requirement

That isn't because employers think they don't matter. Their stated beliefs point the other way. Ninety-three percent agree that nearly every role now carries some technical skill requirement, not just technical roles.

Eighty-seven percent agree that underlying skill and judgment matter more than how fluently someone uses AI tools. Nine in ten agree a candidate can be fluent with AI and still be a weak hire. This is the clearest gap between belief and practice in the study, and it's widest exactly where there's no track record to lean on. (See Figure 2)

Even at Box, an AI-first company, the expectation hasn't reached hiring outside engineering yet.



"I don't know if they're quite as far along in terms of adding AI fluency, at least a technical assessment, to those non-technical hiring processes."

— Megan Yip, University Recruiting, Box

At Lyft, most early-career hiring is technical. The business roles that come through are one-offs: marketing one year, then global supply chain, real estate development, account management.



"A lot of these roles we haven't done before. We have new roles every year."

— Vivian Chen, Principal Recruiter, Lyft

Seven in ten say they assess AI fluency. A third take the candidate's word for it.

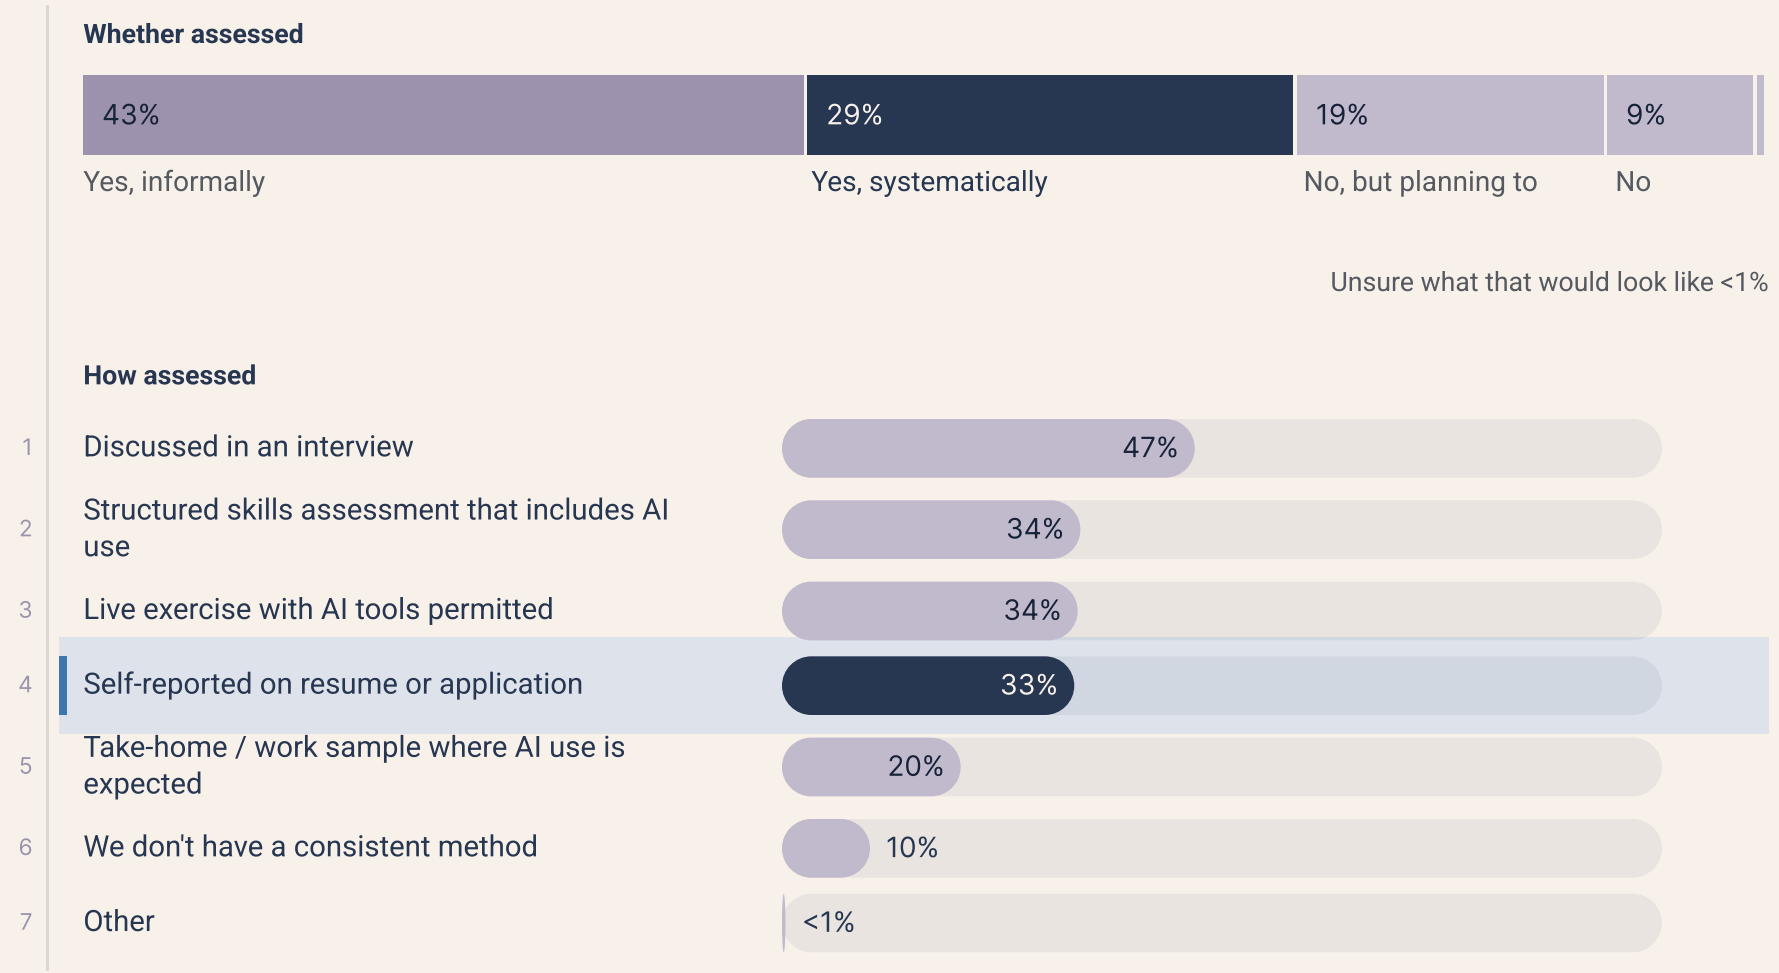
AI fluency at least gets attention. More than seven in ten employers say they assess it. But fewer than three in ten do it systematically.

The rest do it informally, and a third of all employers simply take the candidate's word for it, from a self-reported line on a resume or application. That's the same document this report has already shown to be losing signal.

Judgment has no method at all. There's no instrument in the survey that employers point to, and no pattern in the data that suggests one exists.

Figure 11

How employers assess AI fluency



Share of all employers. How assessed is multi-select.

The people closest to this problem don't have an answer either

Asked what separates candidates with good AI judgment from the rest, a principal recruiter running technical early-career hiring at a major technology company answered without hedging:

> "I don't know. What is the answer to that? When you were asking, I was like, oh my gosh, I want to know what the answer is."

The reason the question matters more now than it did two years ago is that the entry-level job itself has moved.



> "More of the work that an entry-level person is doing is done by AI now. In the past they might have written the code and somebody else reviews it. Now the code is being written by AI, so their job is to review. It's almost like your role has expanded, or shifted upwards."
>
> — Megan Yip, University Recruiting, Box

Fluency may be the on-ramp to judgment

There is one hopeful thread. Employers who have built a way to assess AI fluency are the ones who start noticing judgment underneath it. Watching someone use AI to solve a problem shows you whether they catch its mistakes.

Fluency, in other words, may be the on-ramp to judgment.

> "AI is not always going to be 100% correct. It can do a great leg of the work, but if it's not exactly what you're looking for, are you able to then correct it to fit the solution you need?" says Yip.

At Box, the strongest early signal came from a debugging exercise in a language most students haven't learned, with AI assistance switched on. Candidates who could use the tool and catch what it got wrong separated from those who couldn't.



> "That one actually had the most variance, and it turned out to be probably the truest signal we had at the top of the funnel."
>
> — Megan Yip, University Recruiting, Box

Databricks took a similar route. Instead of asking candidates to write code, it hands them a large repository they haven't seen and asks what's wrong with it.



> "It's almost like doing more of a quality check on what's wrong with it, why, what would you do?"
>
> — Ashley Santos, Databricks

And a second approach is emerging: Reading what the candidate asked the AI.

Some assessment tools now log the candidate's prompts, which turns an unobserved exercise into something closer to observable.



"Without seeing what they're actually asking the chat agent, I don't know how else we would see what kind of questions they're asking. From seeing enough of what people are asking, then we can stratify into this is good use of AI that shows judgment, versus somebody who keeps hacking at it trying to get it to solve."



— Vivian Chen, Principal Recruiter, Lyft

Section 2:

Where the Signal Comes Back



- Finding 5: The resume didn't lose its signal because candidates started using AI. It lost signal because we never set the rules.

The resume isn't dead. It went dark, and the employers who set rules for candidates are still reading it fine.

The resume is still the first tool. Only two in five get much from it.

Start with how the market treats the resume. It's still the most-used screening method by ten points, and nine in ten employers say it gives them at least some signal.

But only about two in five say it gives them a lot of signal. Something is separating the employers who can still read the resume from the ones who can't.

Some employers have stopped asking the resume to carry much weight at all.

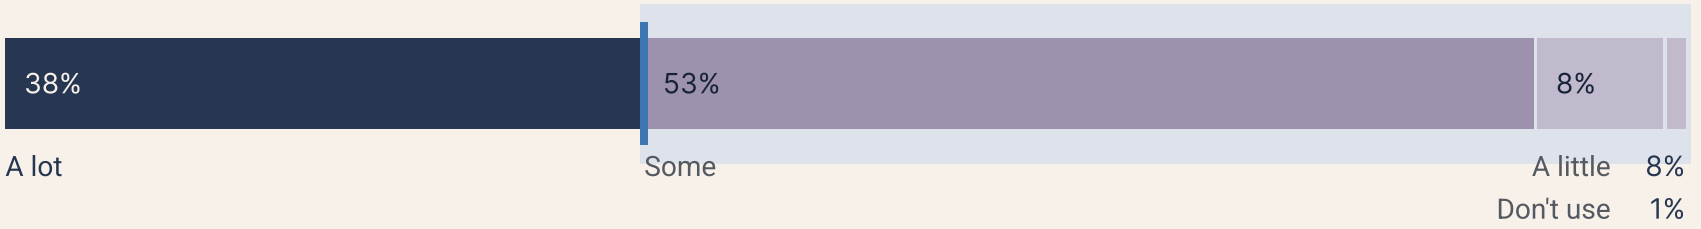


"We take it very lightly. We try not to overindex on it."

— Ashley Santos, Databricks

Figure 12

How much signal the resume gives



How much signal the resume gives about early-career candidates. Share of all employers.

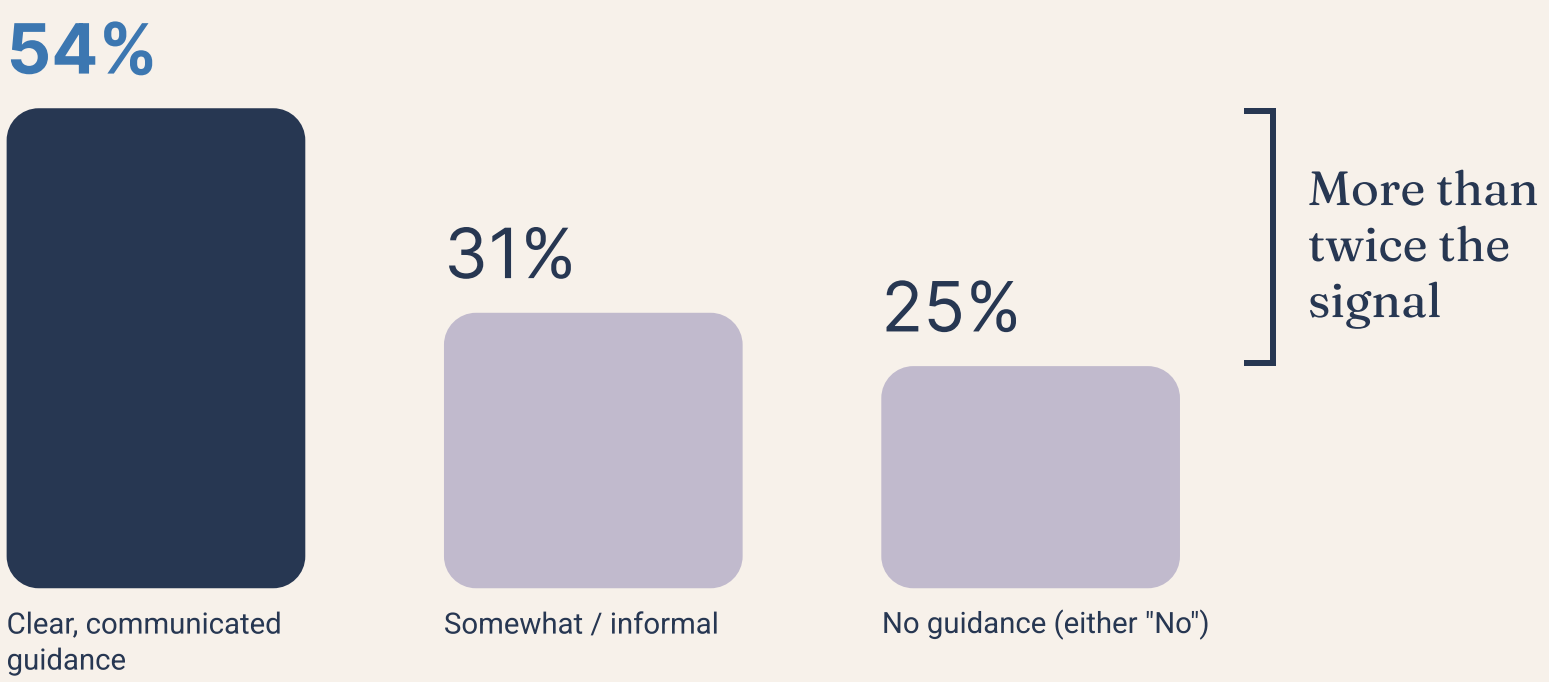
Clear rules more than double the signal

That something is guidance. Among employers who give candidates clear, communicated rules on appropriate AI use, the resume delivers strong signal 54% of the time. Among employers who give no guidance, it's 25%. Same document, same AI tools available to every candidate, more than twice the signal.

The effect is specific to the resume. It holds after accounting for how worried an employer is about candidates gaming the process. And it makes intuitive sense. When nobody knows what's allowed, a polished resume could mean anything: a strong candidate, a strong prompt, or both. When the rules are clear, employers can read the document for what it is, and candidates who follow the rules become distinguishable from those who don't.

Figure 13

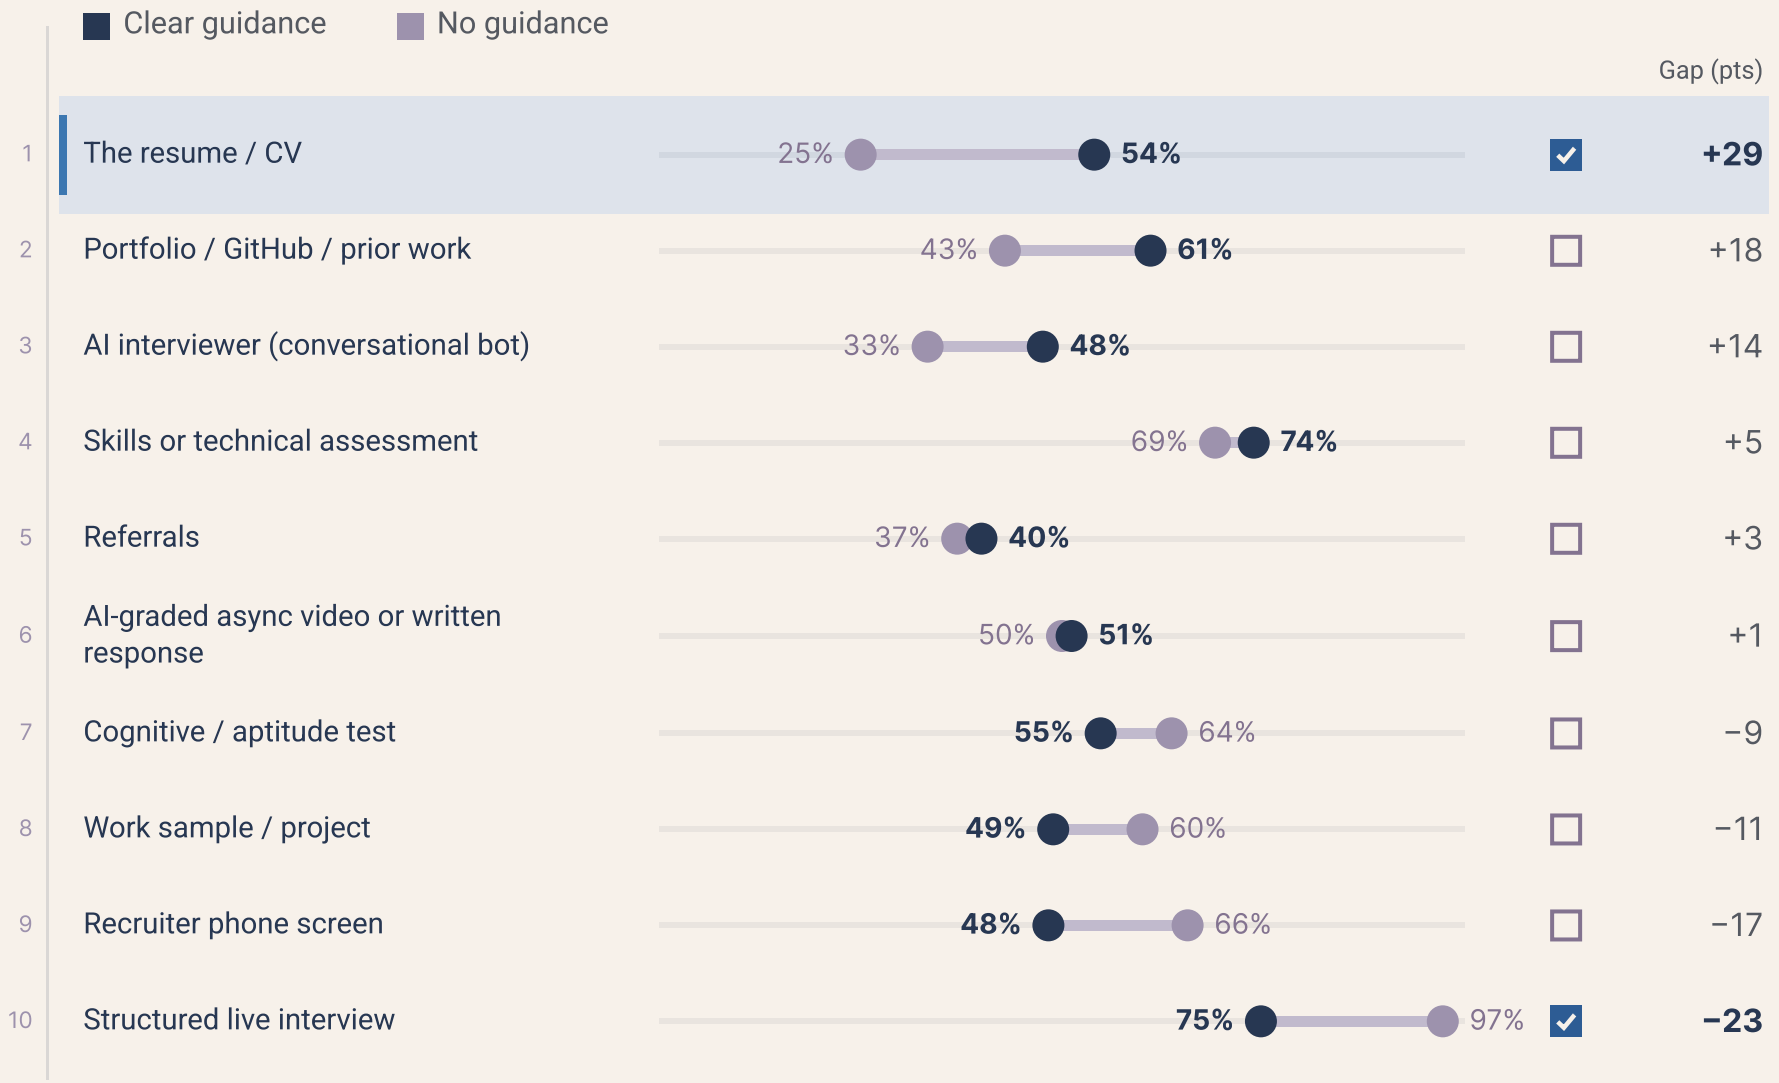
Resume signal by candidate AI-use guidance



Resume gives “a lot” of signal. Base: employers who use the resume. Clear vs no guidance: Fisher p = .002.

Figure 14

Guidance effect is specific to the resume



Share of each method's users who say it gives "a lot" of signal, by candidate AI-use guidance. Base: users of each method. Filled check: gap significant at p < .05 (Fisher exact). Outlined: not significant.

A quarter of employers worry about gaming and have set no rules

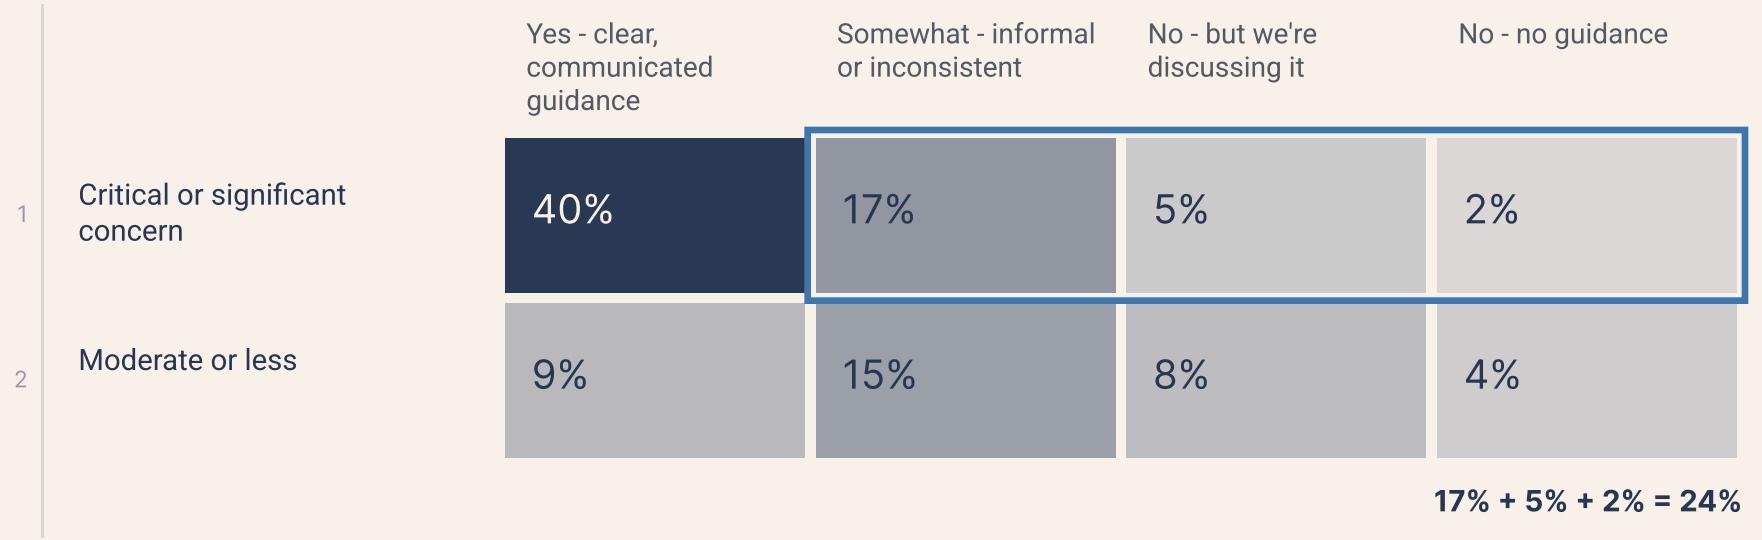
The worst position in the market is common. A quarter of employers are worried about candidates gaming the process and have set no clear rules. They've identified the problem and left in place the condition that causes it.

For most of this market, the concern about gaming is a concern about fraud and identity verification: nearly two in three rate it a critical or significant issue in early-career screening.

That problem deserves its own depth, and it gets it in Kyle & Co's The State of Candidate Fraud Detection and Prevention (June 2026). For this study, the point is narrower. Much of the signal employers blame AI for losing, they could recover by telling candidates the rules.

Figure 15

Concern about gaming × candidate guidance



Cells are shares of the full sample. Outlined: critical or significant concern without clear guidance.

This is the best evidence in the dataset that structure closes the gap AI opened.

Finding 6: Employers who watch candidates work are ahead

The employers furthest ahead on early-career quality of hire aren't asking candidates for more material. They're watching candidates do the work.

Quality of hire climbs with every observed method

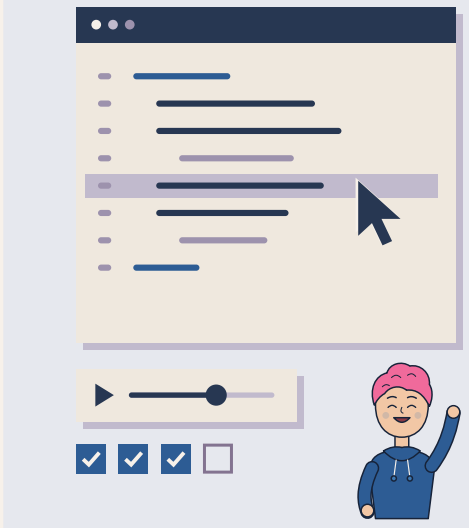
Count the methods where a candidate is observed while doing the work: a skills or technical assessment, or a cognitive or aptitude test. Then look at how employers rate their early-career quality of hire.

Among employers who use neither, 18% say they're ahead of peers. With one, it's 29%. With both, it's 36%, twice the starting rate. Each observed method adds to the share.

Take-homes, portfolios, and async responses show no climb at all

Now count the methods completed away from observation: work samples and take-home projects, portfolios, AI-graded video or written responses.

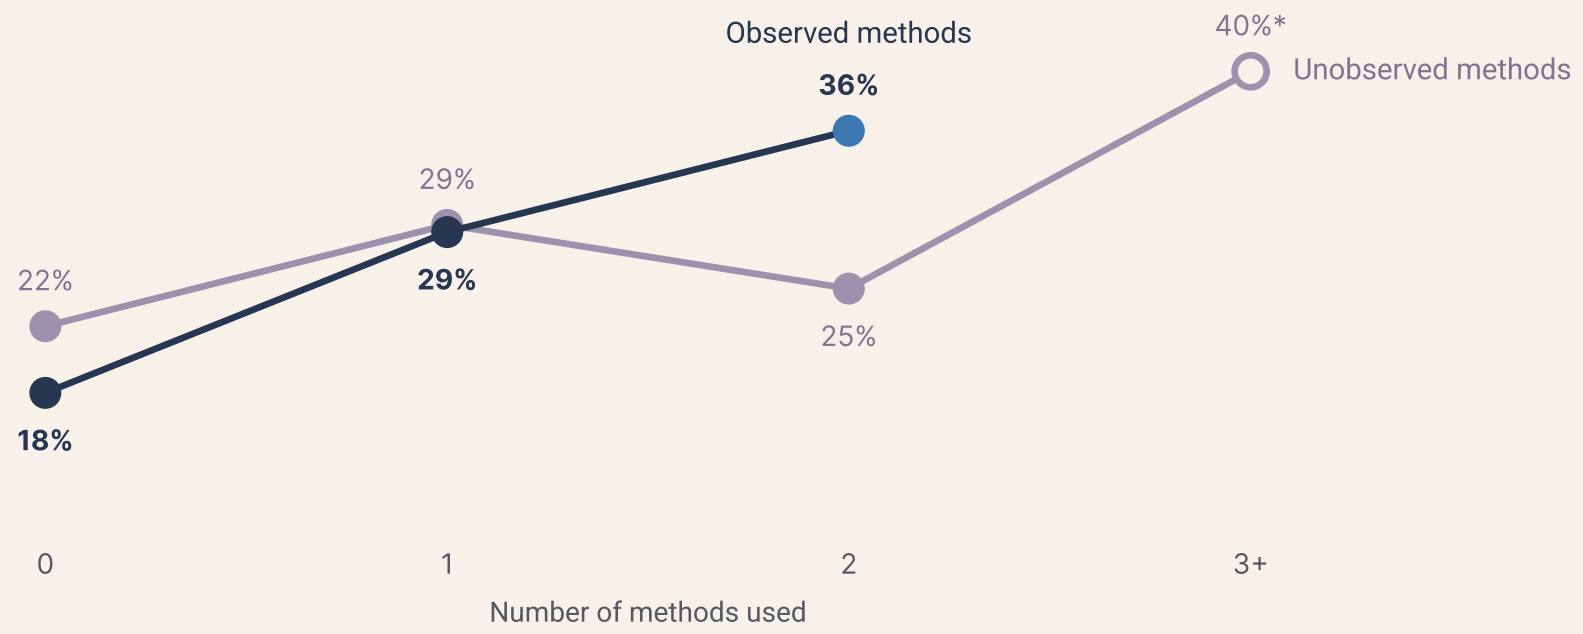
There's no climb at all. Employers who use more of them rate their quality of hire about the same as employers who use fewer. The candidate still produces something. What's missing is anyone watching it get made.



> “
>
> "Almost everyone's going to have a pretty high score on the take-home portion. The live interview is definitely the biggest indicator of how well they really do,” says Yip.

Figure 16

Quality of hire climbs with observed methods, not unobserved ones



Share rating early-career quality of hire “Ahead of peers.” Trend p (Cochran-Armitage): observed .018, unobserved .367.
* Unobserved 3+ bucket is about 2% of the sample. Read with caution. Observed has two methods, so no 3+ point.

An observed exercise is where the evidence and the person stay connected

That distinction matters more now than it did five years ago. A take-home project or a written response could always be polished with outside help. With AI, it can be produced outright.

An observed exercise is one of the few places left where the evidence and the person are guaranteed to be connected. It's also where employers get to see the raw materials this report is about: how someone approaches a problem they haven't seen, what they do when they get stuck, and whether they notice when a tool gives them a wrong answer.

This is an association–not proof–but the pattern is consistent, it runs in one direction, and it separates cleanly by whether anyone was watching. It also lines up with what leaders do, which is where this report goes next.

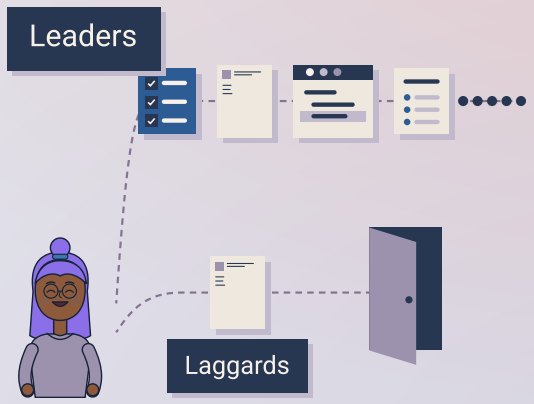
Employers aren't ignoring the problem with unobserved work. Some are trying to make it observable. At Lyft, the first time anyone sees an early-career candidate's work is a take-home technical assessment, and the open question is how to see the judgment inside it.



> "Even if it's a take-home and they can use AI, when they present, they still need to know what they've put together."
>
> — Vivian Chen, Principal Recruiter, Lyft

Section 3:

How Leaders Differ from Laggards



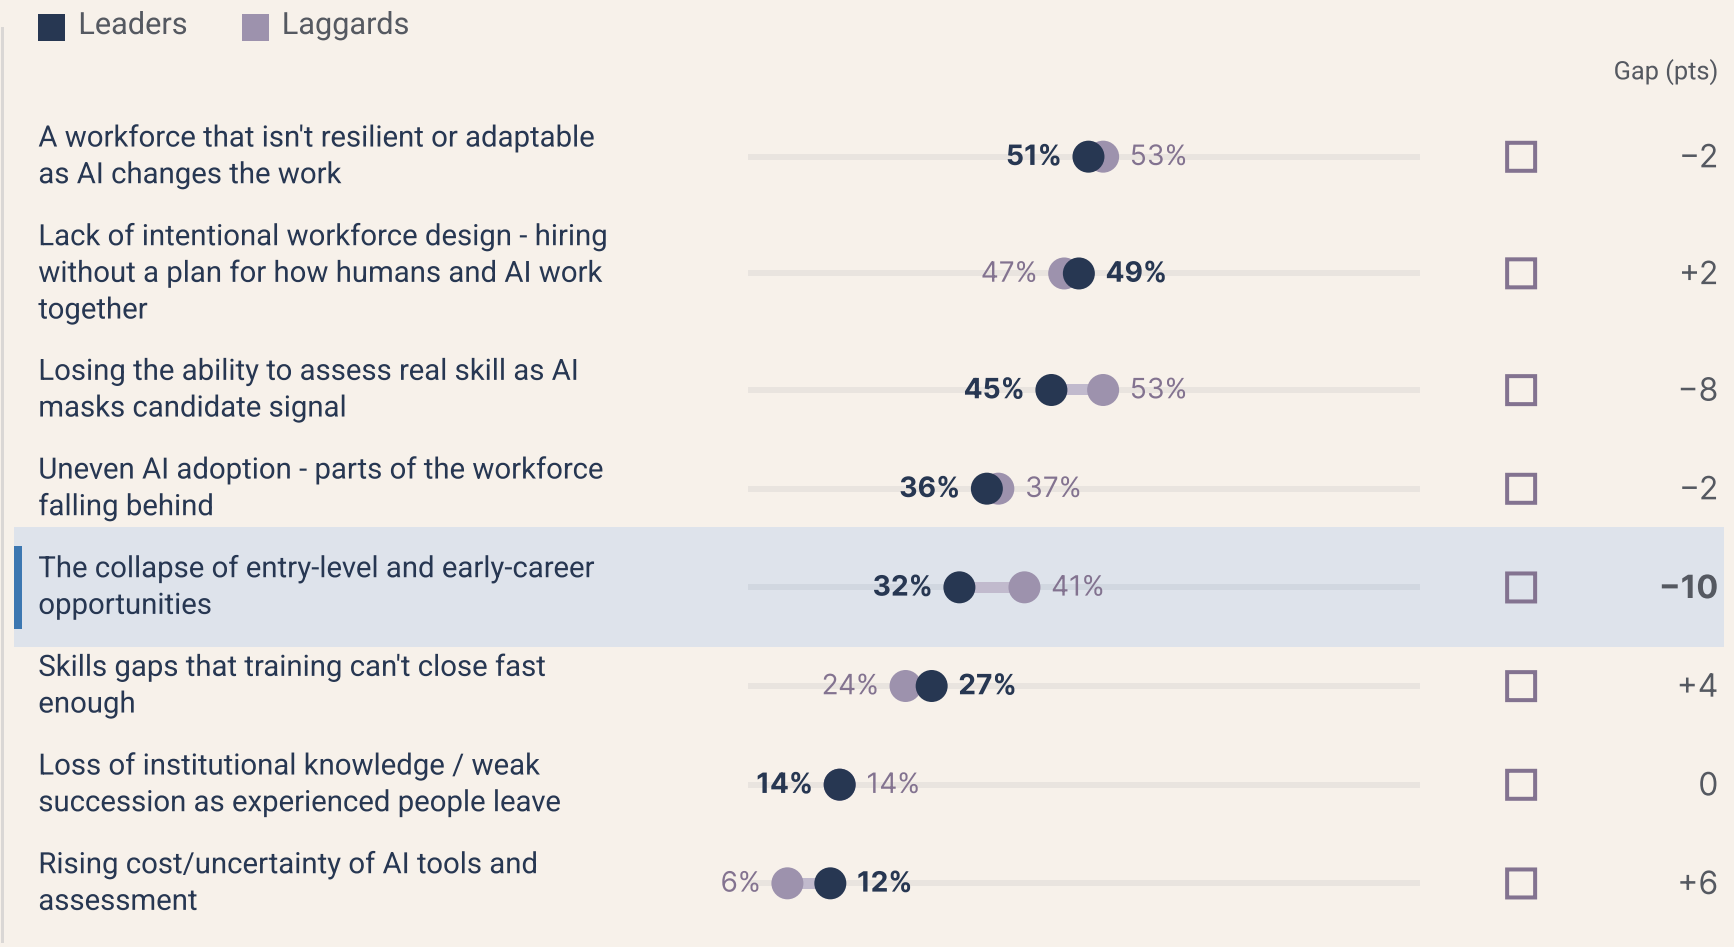
- Leaders haven't out-thought anybody. They've out-built them.

The obvious assumption is that the employers doing best at early-career hiring understand something the rest don't. The data says otherwise.

Leaders and laggards worry about the same risks in roughly the same proportions. Half of each group names workforce resilience. Their concern about workforce design, uneven AI adoption, and skills gaps is nearly identical. Leaders aren't seeing a different problem.

Figure 17

Leaders and laggards worry about the same things

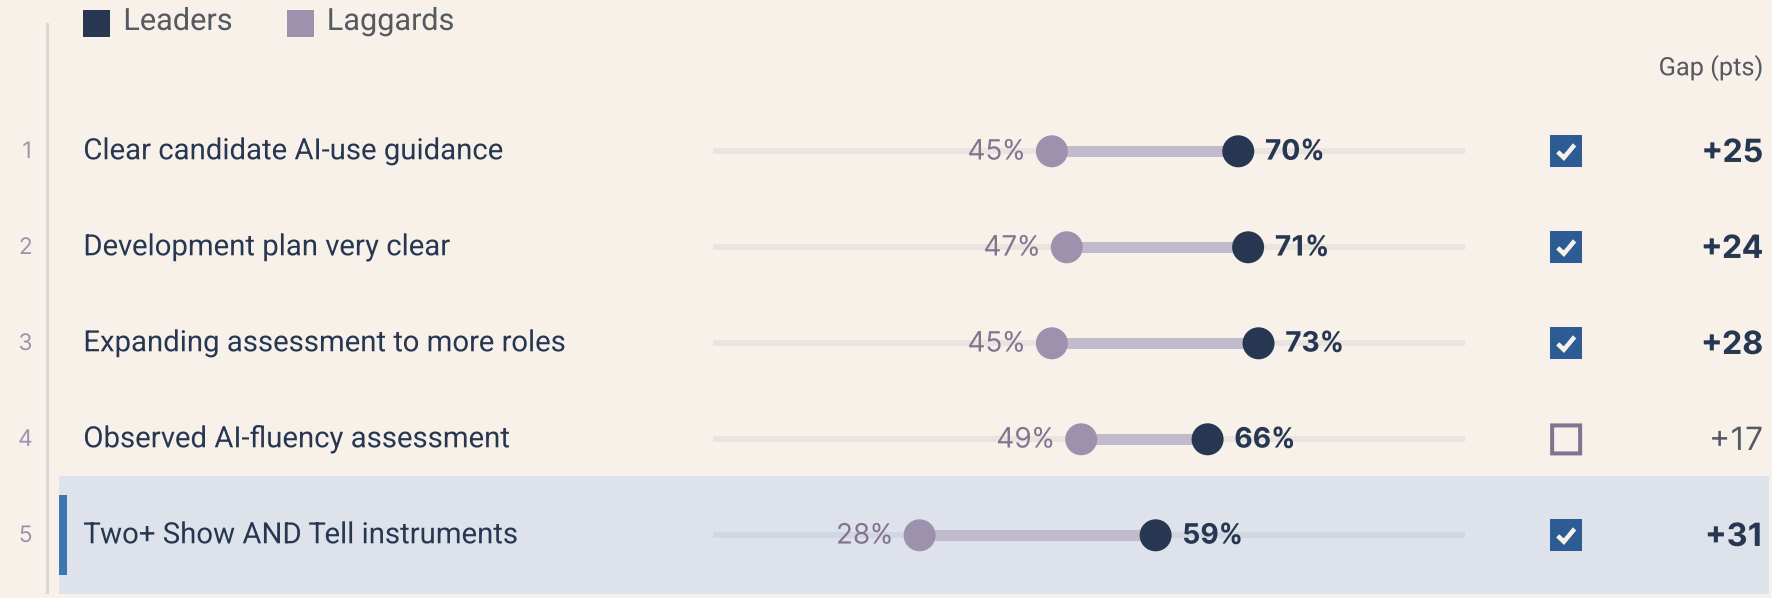


Share of leaders and laggards naming each risk in their top three.
Filled check: gap significant at p < .05 (Fisher exact). Outlined: not significant.

What's different is what they've put in place. Across the survey, five practices separate leaders from laggards. None of them is exotic. Most are the kind of thing a talent team could start next quarter. Together, they describe an organization that has decided to look at candidates directly and to tell them the rules while doing it.

Figure 18

Practice gaps at a glance



Share of leaders and laggards using each practice.
Filled check: gap significant at p < .05 (Fisher exact). Outlined: not significant.

AI momentum built the rules. It didn't change how employers see candidates.

If leaders aren't thinking differently, maybe they're just further along on AI. They're not. Using the convention from Kyle & Co's AI Momentum Model research, we compared organizations accelerating fastest on AI (those describing strategic acceleration or an integrated approach over the past year) with those just starting out or pulling back.

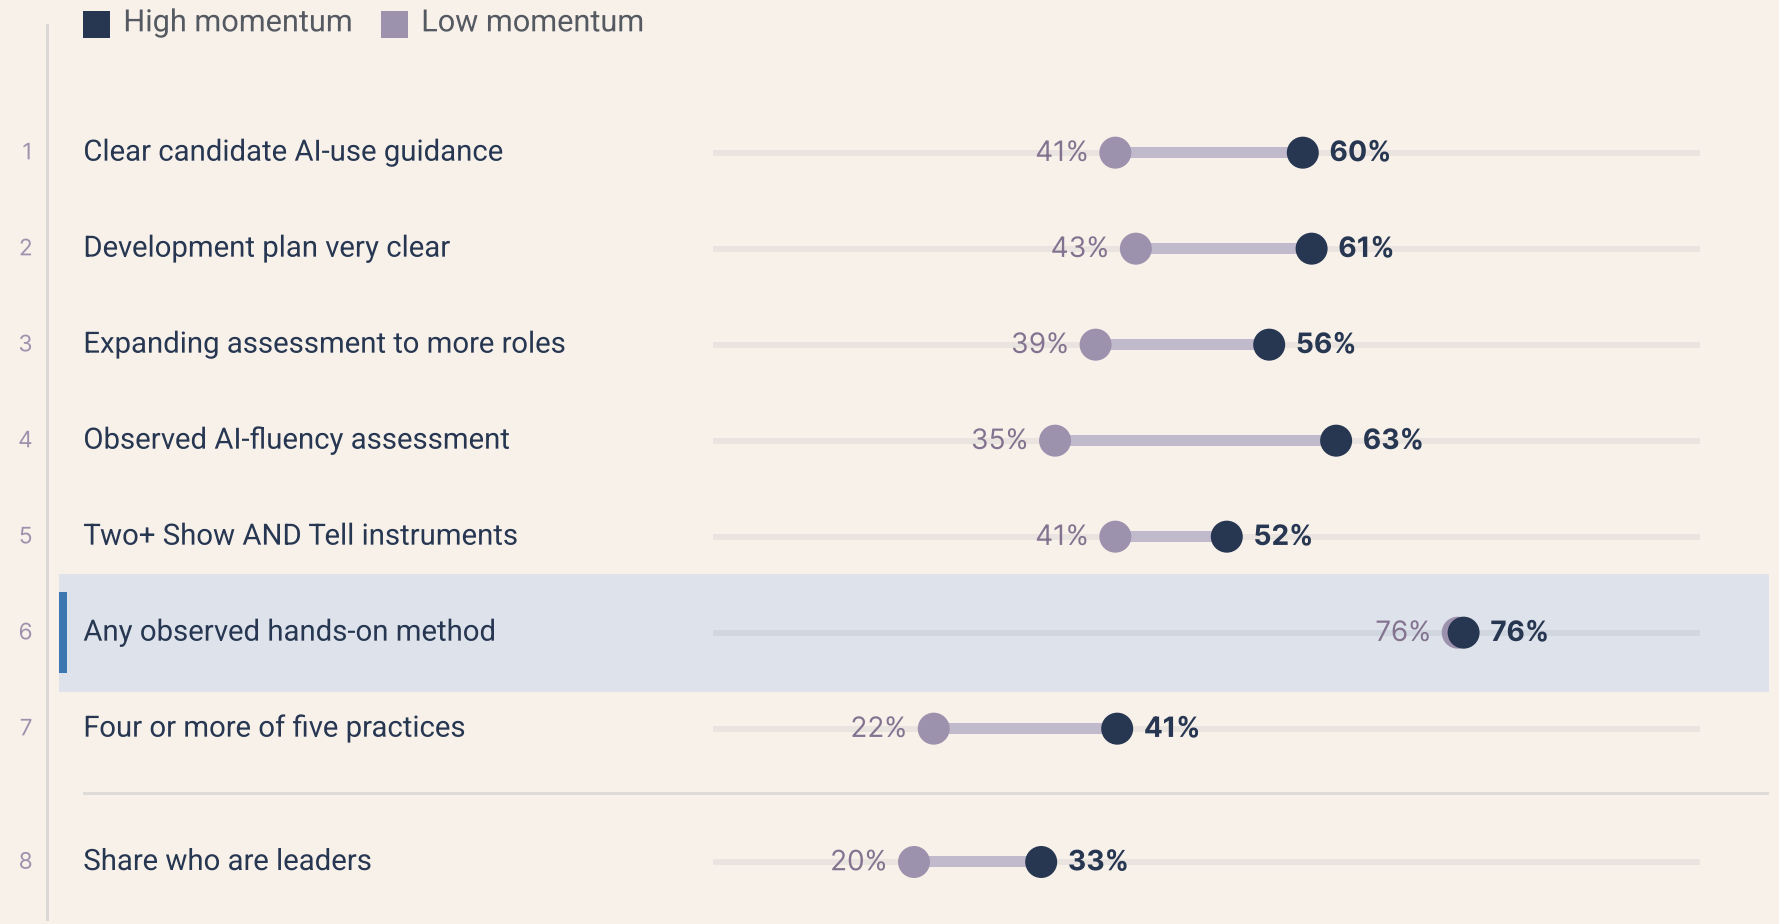
The fast movers are only modestly more likely to be leaders, and nearly half of all leaders aren't among them.

AI-forward organizations do share some leader habits. They're more likely to give candidates clear rules on AI use, to have a development plan someone can describe, and to be extending assessment to new roles. That's governance, and they've built it.

What they haven't changed is how they see candidates. AI-forward organizations use observed assessment methods at exactly the same rate as organizations that are just beginning. Momentum built the rules. It didn't change the evaluation, and the evaluation is what tracks with outcomes.

Figure 19

What AI momentum does and doesn't come with



High = strategic acceleration or integrated (43% of sample); low = exploration or paused (18%). Share of each group with each practice.

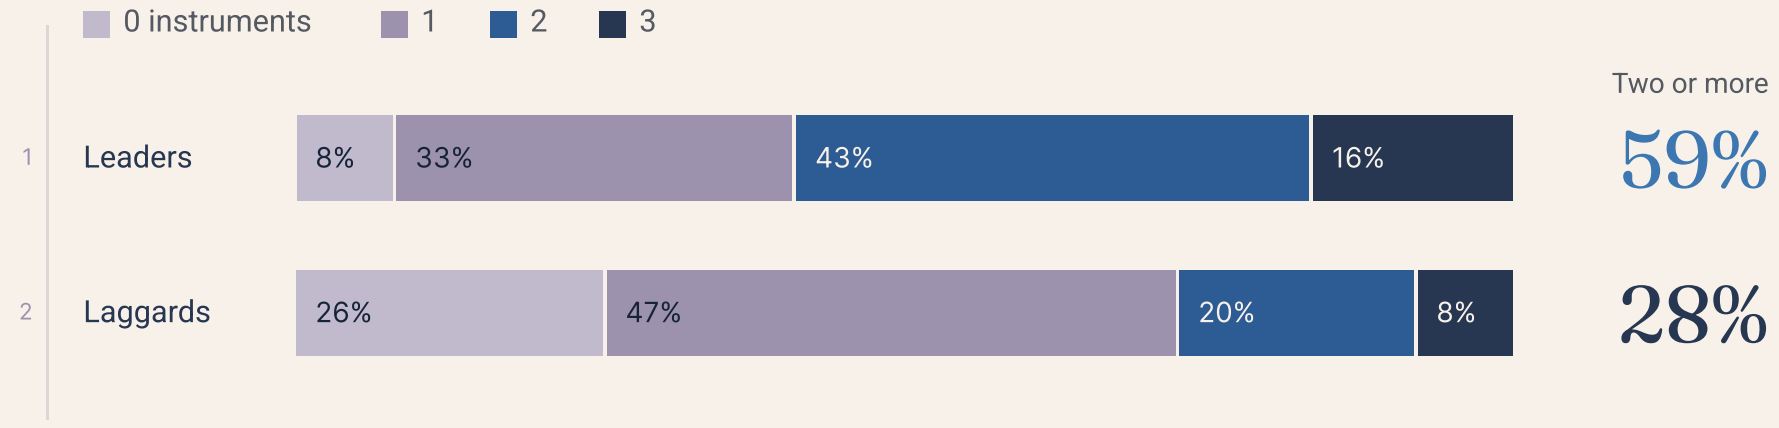
1. Show AND Tell: Leaders kept the resume and added evidence around it

Leaders didn't drop the resume. They use it at the same rate as everyone else. What they added is a second kind of evidence around it: evidence beyond the document.

We looked at three sources of evidence beyond the document: a skills or technical assessment, a structured live interview, and a portfolio of prior work. Nearly three in five leaders use at least two of them before making an early-career hire. Among laggards, it's barely more than a quarter. And a quarter of laggards use none of the three.

Figure 20

Show AND Tell



Number of Show AND Tell instruments used: skills/technical assessment, structured live interview, portfolio. Resume use: leaders 73%, laggards 71%.

This is the largest gap between leaders and laggards on any single practice in the study. The emphasis on AND is deliberate. Leaders haven't replaced telling with showing. They ask for both, and they let each check the other.

2. Widen the perimeter: Leaders know the gap isn't in engineering, it's at the frontline

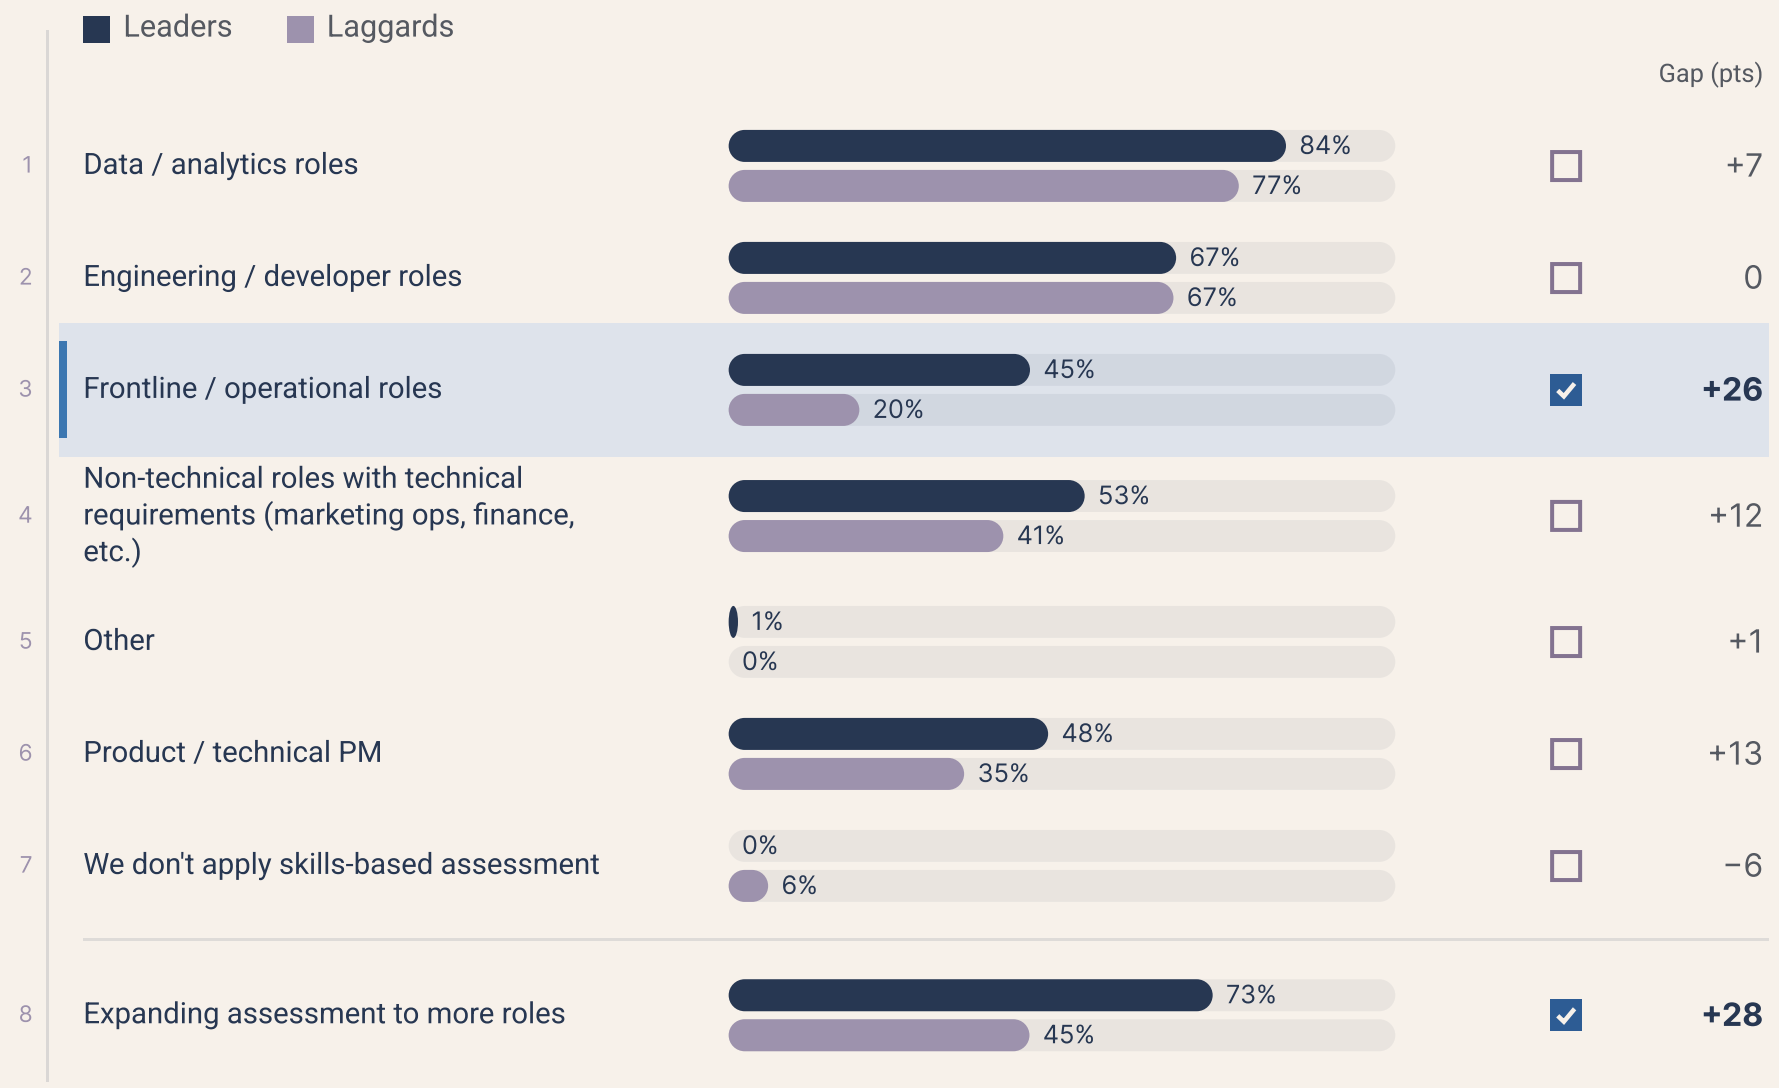
Nearly three-quarters of leaders are extending skills-based assessment to more roles. For laggards, it's closer to a coin flip.

The difference isn't in engineering. Both groups assess engineering and developer roles at identical rates. The gap opens further from the core of technical hiring.

Leaders assess frontline and operational roles at more than twice the laggard rate. They've taken the view that if nearly every role now carries a technical requirement, and most employers agree it does, then nearly every role deserves a real look at whether the candidate can meet it.

Figure 21

Assessment coverage by role family



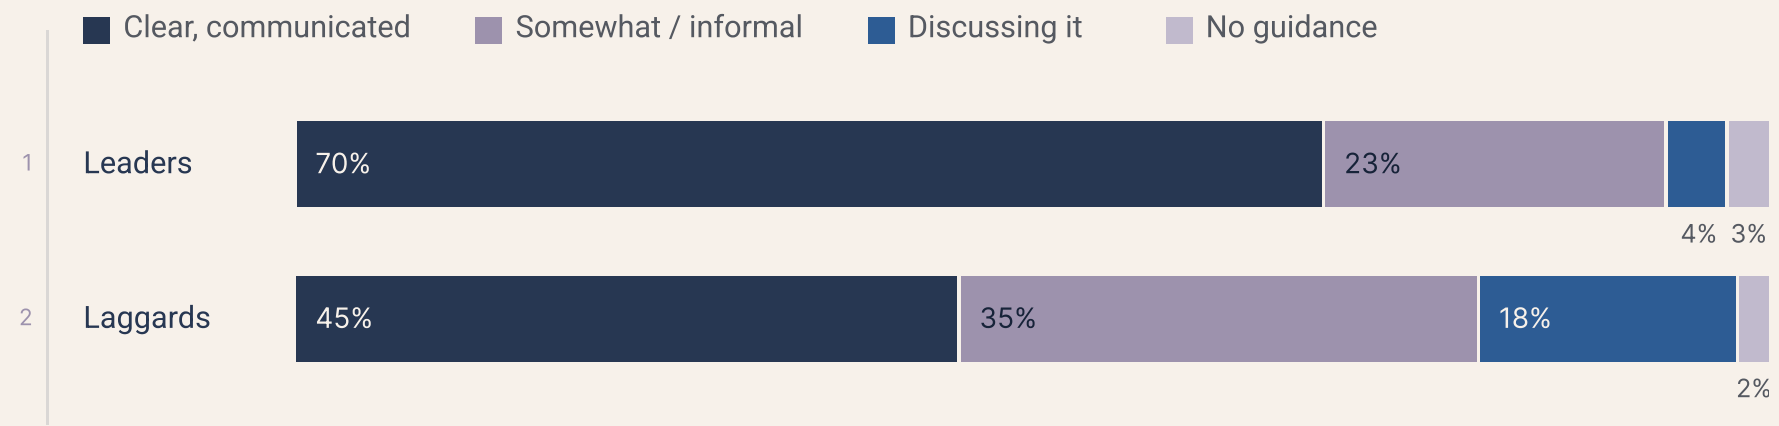
Share of leaders and laggards applying skills-based assessment to each role family.
Filled check: gap significant at p < .05 (Fisher exact). Outlined: not significant.

3. Write the Rules: Seven in ten leaders tell candidates what's allowed

Seven in ten leaders give candidates clear, communicated guidance on appropriate AI use. Among laggards, it's 45%. Laggards are four times as likely to say they're "still discussing it."

Figure 22

Candidate AI-use guidance



Does your organization give candidates guidance on appropriate AI use? Share of leaders and laggards.

This is the practice behind finding #5. The resume holds its signal where candidates know the rules, and leaders are the ones who've told them. It's also the cheapest practice on this list. Writing the rules costs a paragraph in the application.

The rules also decide where the signal shows up. Box changed its AI rules three years running: first no rules, then a ban candidates ignored, then AI allowed on the take-home and switched off in the live round. With AI off in the live round, fewer candidates reached the top rating.



> "It's shifted the true signal quite a bit further down the funnel."
>
> — Megan Yip, University Recruiting, Box

At Lyft, the concern was that application responses had begun to sound alike.



> "I felt like I kept getting the same answer. Lyft is at the intersection of these three things. This is like the 20th response that said it's at the intersection of these three things. I thought, 'This is not helpful.'"
>
> — Vivian Chen, Principal Recruiter, Lyft

Lyft added guidance beside the question explaining why recruiters wanted candidates' own words.

> "I said: we send and read thousands of these responses, so we're able to tell right away if the answer is AI generated, because certain phrases and concepts come up repeatedly. What we're most interested in is your own genuine voice."

Chen reports seeing more personalized stories since adding the guidance.

> "At least the answers feel more real to me. I see a lot more personalized stories, which is exactly what I want."
>
> — Vivian Chen, Principal Recruiter, Lyft

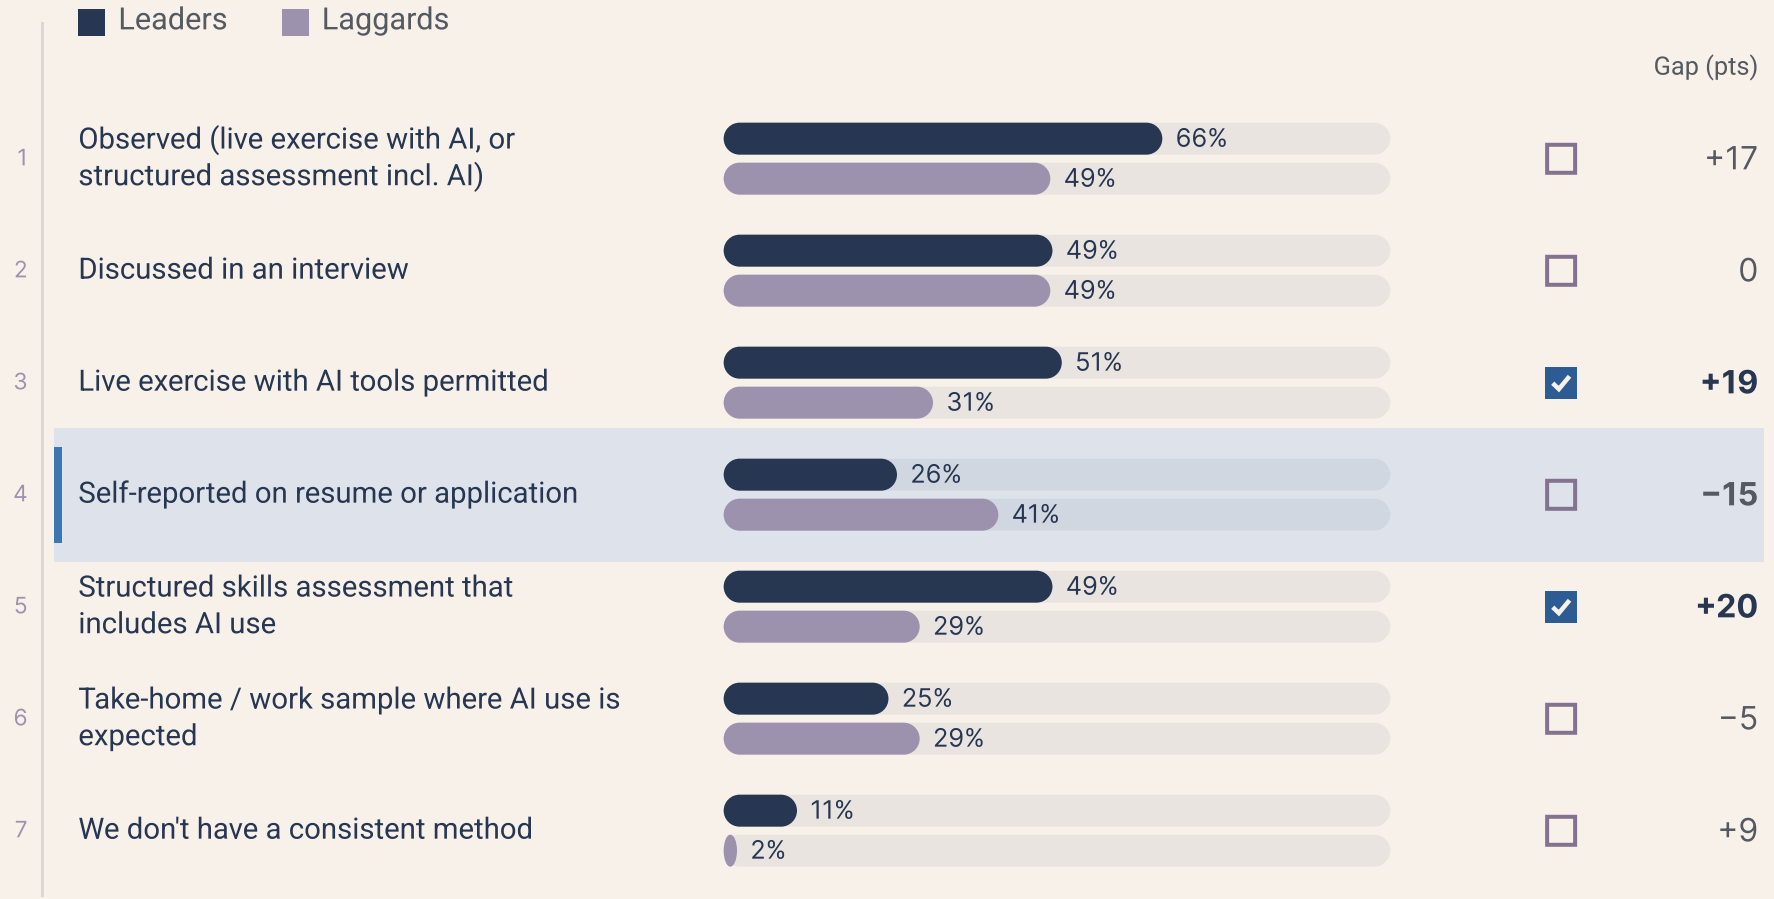
4. Watch instead of asking: Leaders observe AI fluency, laggards take it on faith

Both groups talk about AI fluency in interviews at the same rate, about half. That's where the similarity ends.

Leaders are 17 points more likely to actually observe AI fluency, through a live exercise where AI use is permitted or a structured assessment that includes it. Laggards are more likely to rely on the candidate's own claim: 41% take a self-reported line on a resume or application as evidence, against 26% of leaders.

Figure 23

How leaders and laggards assess AI fluency



Share using each method to assess AI fluency. Base: full cohort, including employers who don't assess AI fluency.
Filled check: gap significant at p < .05 (Fisher exact). Outlined: not significant.

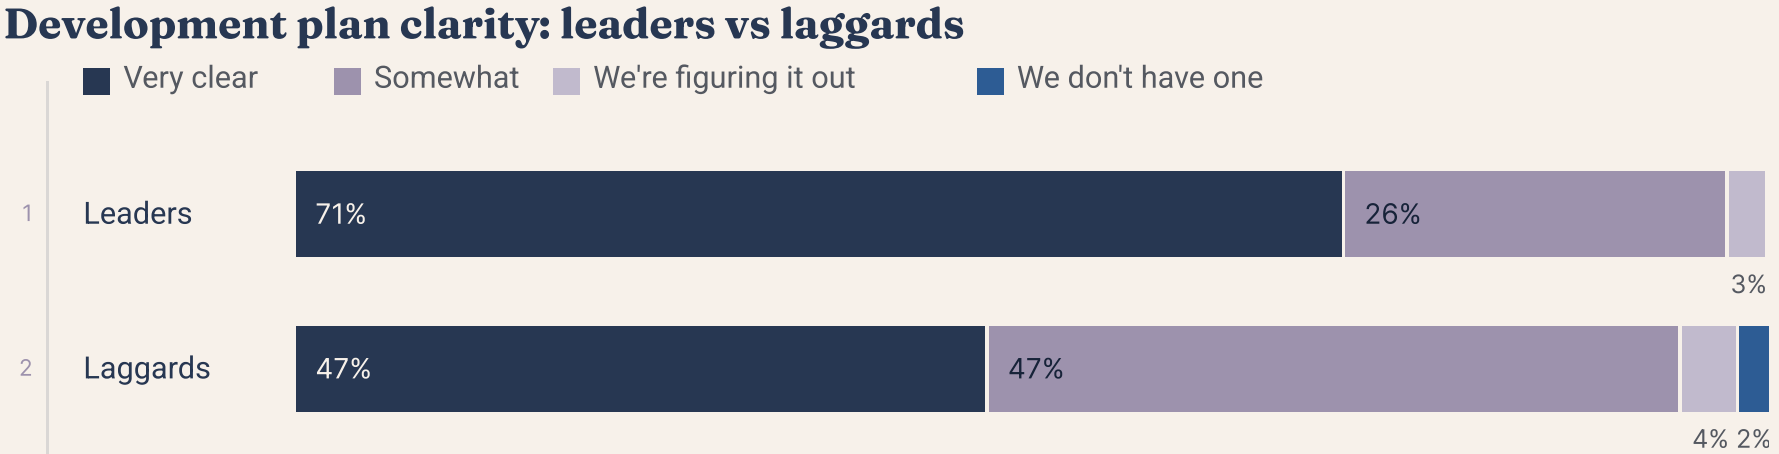
The pattern mirrors finding #6. The evidence that separates leaders is the evidence someone watched happen.

5. Have a plan someone can describe–the best performers all do

Seven in ten leaders rate their plan for developing early-career hires as very clear. Fewer than half of laggards say the same.

Figure 24

Development plan clarity: leaders vs laggards

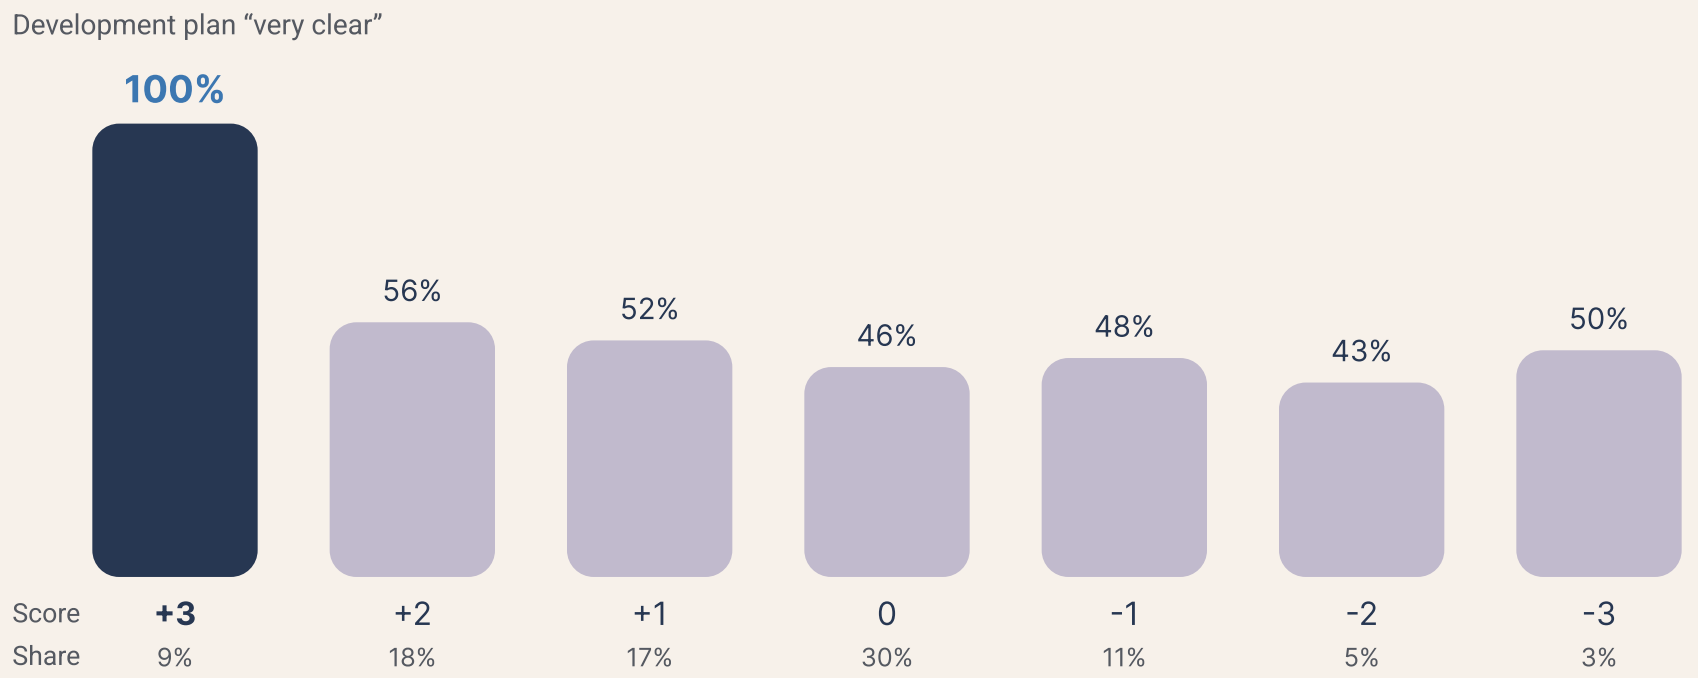


How clear is the development plan for early-career hires? Share of leaders and laggards.

The strongest version of this finding sits at the very top of the market. Among employers who rate themselves ahead of peers on all three early-career measures, about one in eleven, every single one describes their development plan as very clear. It's the one practice in the study that every top performer shares.

Figure 25

Development plan clarity by outcome score



Outcome score: ahead (+) or behind (−) on the three early-career KPIs. Share = share of employers measuring all three at each score.

That matters because hiring early-career talent is a bet on raw materials. The development plan is how the bet pays off.

At Databricks, AI competencies now run from hiring into performance reviews, with different expectations at each career level. For early-career hires, the bar is about how they learn.



"For the real junior folks, what we really try to assess for is: where is their curiosity, where's their builder mindset, are they expressing interest, are they looking into different trainings?"

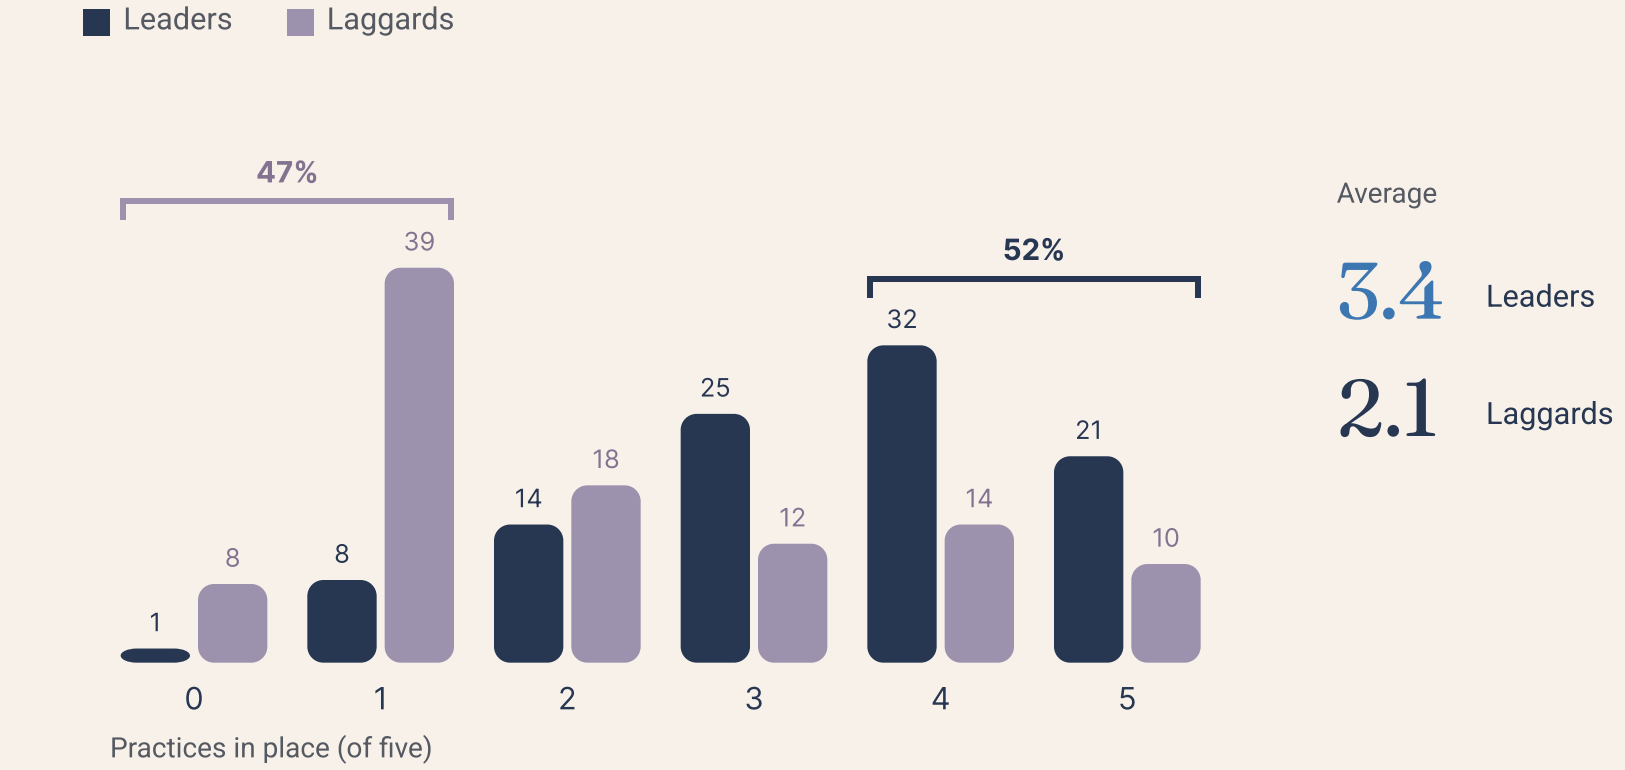
— Ashley Santos, Databricks

Best practices stack: The distance between the top and the bottom is about one practice

These five practices travel together. More than half of leaders run four or more of them. Nearly half of laggards run one or none.

Figure 26

The stack



Share of leaders and laggards by number of the five practices in place. Average is a count of practices.

The average difference between the two groups is about one practice. That's the most encouraging number in this report. The distance between the top and the bottom of the market isn't a transformation program. It's roughly one thing done well.

Leaders claim less: the best performers are the most willing to say what's unfinished

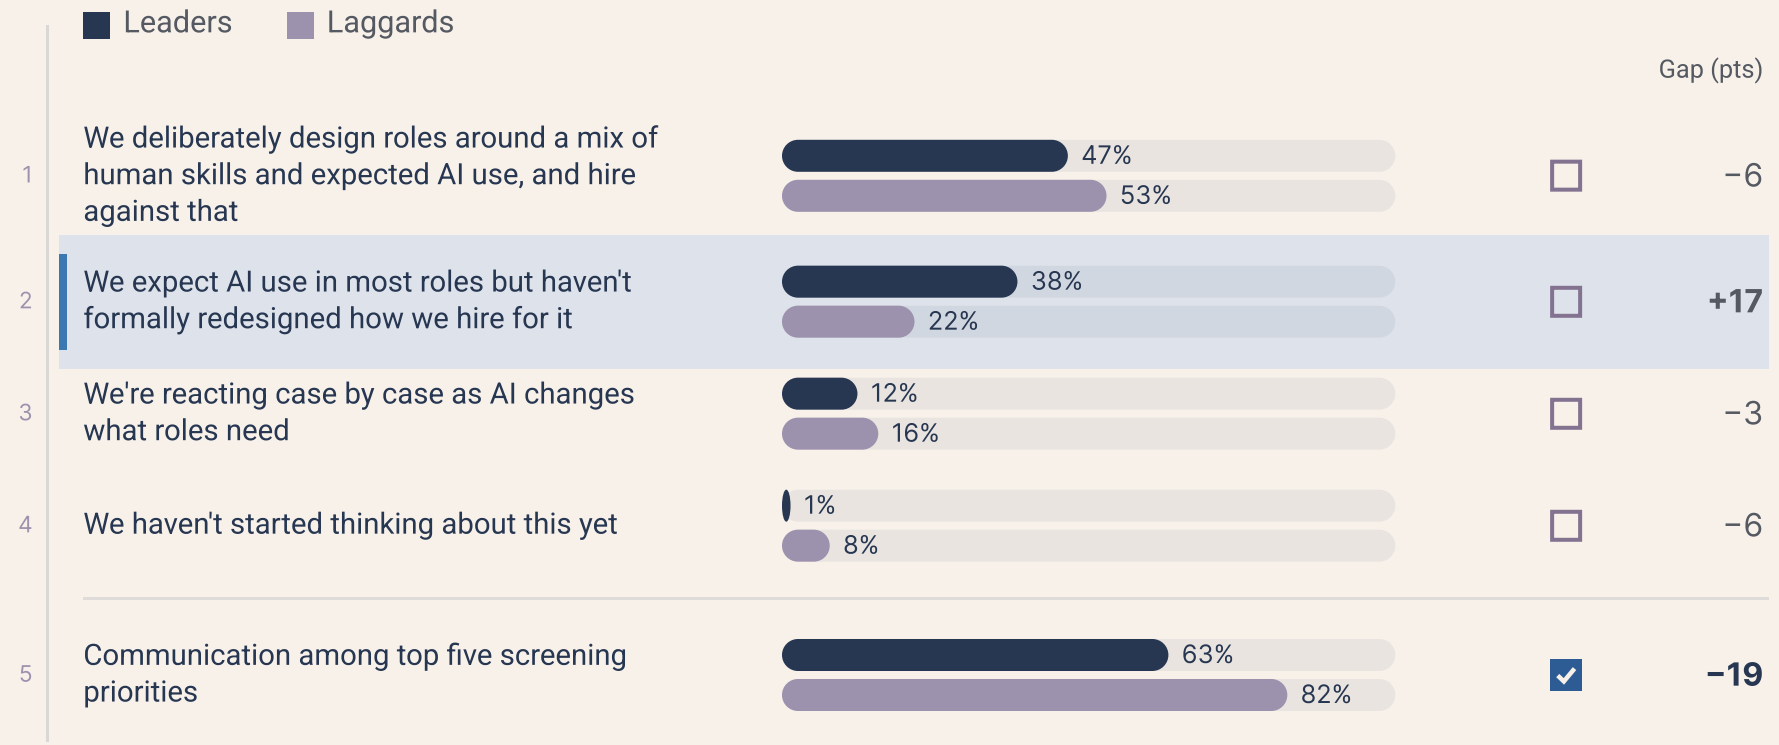
One last pattern runs against expectation. The employers getting the best results are the most willing to say what's unfinished.

Leaders are nearly 20% more likely than laggards to say they expect AI use in most roles but haven't formally redesigned how they hire for it.

Comparatively, laggards lean toward saying they've already deliberately designed roles around AI. Laggards also over-index on communication as a top screening priority, the quality most easily read from a conversation or a document.

Figure 27

Leaders claim less



Share of leaders and laggards choosing each statement on hiring for AI use. Communication row: share naming it among their top five screening priorities.
Filled check: gap significant at p < .05 (Fisher exact). Outlined: not significant.

The group with the strongest outcomes is the one holding itself to the harder standard. It's a useful reminder for any reader benchmarking against this report: confidence about having solved the problem is not evidence of having solved it.

Close the gap



The gap between leaders and laggards is about one practice wide. These are the moves that close it, roughly in order of cost.

1

Write the rules, and put them in the application itself

Tell candidates what AI use is permitted, and put it in the application itself, not on a careers page they may never see. Be specific: what's fine for the resume, what's fine for a take-home, what's expected in a live exercise.

This is the cheapest change in the report, and it's the one most directly tied to recovering the resume's signal.

2

Add one observed instrument before you make an offer

If no one currently watches an early-career candidate do the work before you make an offer, start there. Going from none to one observed method is the largest single step in the quality-of-hire climb. It doesn't need to be elaborate. A short live exercise on a realistic problem shows you more about someone's raw materials than another round of documents.

3

Watch AI use instead of asking about it

Stop taking self-reported AI fluency at face value. Let candidates use AI in a live exercise and watch what they do with it. You'll see fluency, and you'll often see judgment too: whether they check the output, notice when it's wrong, and know when to set it aside.

4

Extend assessment past engineering, starting at the frontline

Most employers already agree nearly every role carries a technical requirement. The leaders are the ones acting on it. Frontline and operational roles are where the gap between leaders and laggards is widest, and where a lightweight, role-relevant assessment is most likely to surface talent that paperwork misses.

5

Make the development plan something a manager can describe

Every top performer in this study has a development plan they'd call very clear. Clear means a hiring manager can say what the first months look like, month by month, and what "on track" means at each point. Hiring for raw materials only pays off if someone has a plan for turning them into skills.

6

Build a way to test judgment, starting from how candidates use AI

Nobody in this study has a settled method, and the practitioners closest to the problem say so plainly. But the first attempts all run through the same door: give the candidate AI, give them a problem the tool gets partly wrong, and watch what they do with the answer.

Box found its truest early signal in a debugging exercise in an unfamiliar language with AI switched on. Some companies are also looking at the candidate's own prompts for evidence of how well they direct the tool.

Neither is a finished instrument. The employers who build one first will have something the rest of the market is looking for.

Our Research Methodology

CoderPad sponsored this study. Kyle & Co designed the survey instrument with input from CoderPad's team, who also reviewed an early read of the results. Kyle & Co ran the analysis, decided what the data could and couldn't support, and wrote every conclusion in this report. The report is Kyle & Co's independent work.

This study surveyed 271 HR and talent acquisition leaders involved in or accountable for early-career hiring, from manager to CHRO, at organizations of 50 to more than 50,000 employees.

Responses were screened for completeness, an attention check, a fictitious-vendor trap, plausibility of reported hiring volume, and straightlining. All figures are shares of the qualified sample or of the group named.

Leaders and laggards. Respondents rated their organization against peers on three early-career outcomes: quality of hire, first-year retention, and speed to productivity. Each rating was scored ahead (+1), on par (0), or behind (-1). Leaders scored +2 or better; laggards scored -1 or worse. Respondents who don't measure one or more of the three outcomes were left unclassified. (See Figure 4.)

Associations, not causes. Every comparison in this report is an association between what employers do and how they rate their own outcomes. Differences were tested for statistical significance. They show where the best-performing employers have made different choices. They don't prove those choices caused the results.

Practitioner interviews. Survey findings are being pressure-tested in interviews with talent leaders, including customers of CoderPad. Kyle & Co conducts every interview, and transcripts are not shared with the sponsor.

About Kyle & Co

Kyle & Co is a modern research and advisory firm helping HR and talent technology leaders make smarter decisions, faster. By blending practitioner insight, rigorous research and strategic advisory, we uncover the signals that matter most, then translate them into action. Our work spans custom research, workforce strategy and market analysis, supporting solution providers and HR leaders alike.

Kyle & Co is the team behind the Human-Centric AI Council, the Transformation Realness podcast and industry-shaping reports on quality of hire, workforce planning and responsible AI. We work with clients like Workday, ADP, GoodTime, Findem and SmartRecruiters.

Contact us: www.kyleandco.com · info@kyleandco.com · linkedin.com/company/kyle-co

CoderPad

About CoderPad

CoderPad helps companies find the right people, pairing hiring process expertise with an AI-powered assessment and live interview platform for engineering, go-to-market, HR, and beyond. From the first screen to the final interview, CoderPad covers the full hiring journey, so teams can see who can do the work.

As AI changes what skill looks like on the job, CoderPad's real-world environment shows how candidates think and work with AI, not just whether they passed a test. The result is faster, clearer, more defensible hiring decisions.

More than 4,000 customers, including two-thirds of the Forbes AI 50, use CoderPad to hold the hiring bar at scale. Candidates prefer it too.