UpGuard

The State of Shadow AI

How employees are using AI at work, and what it means for security leaders.

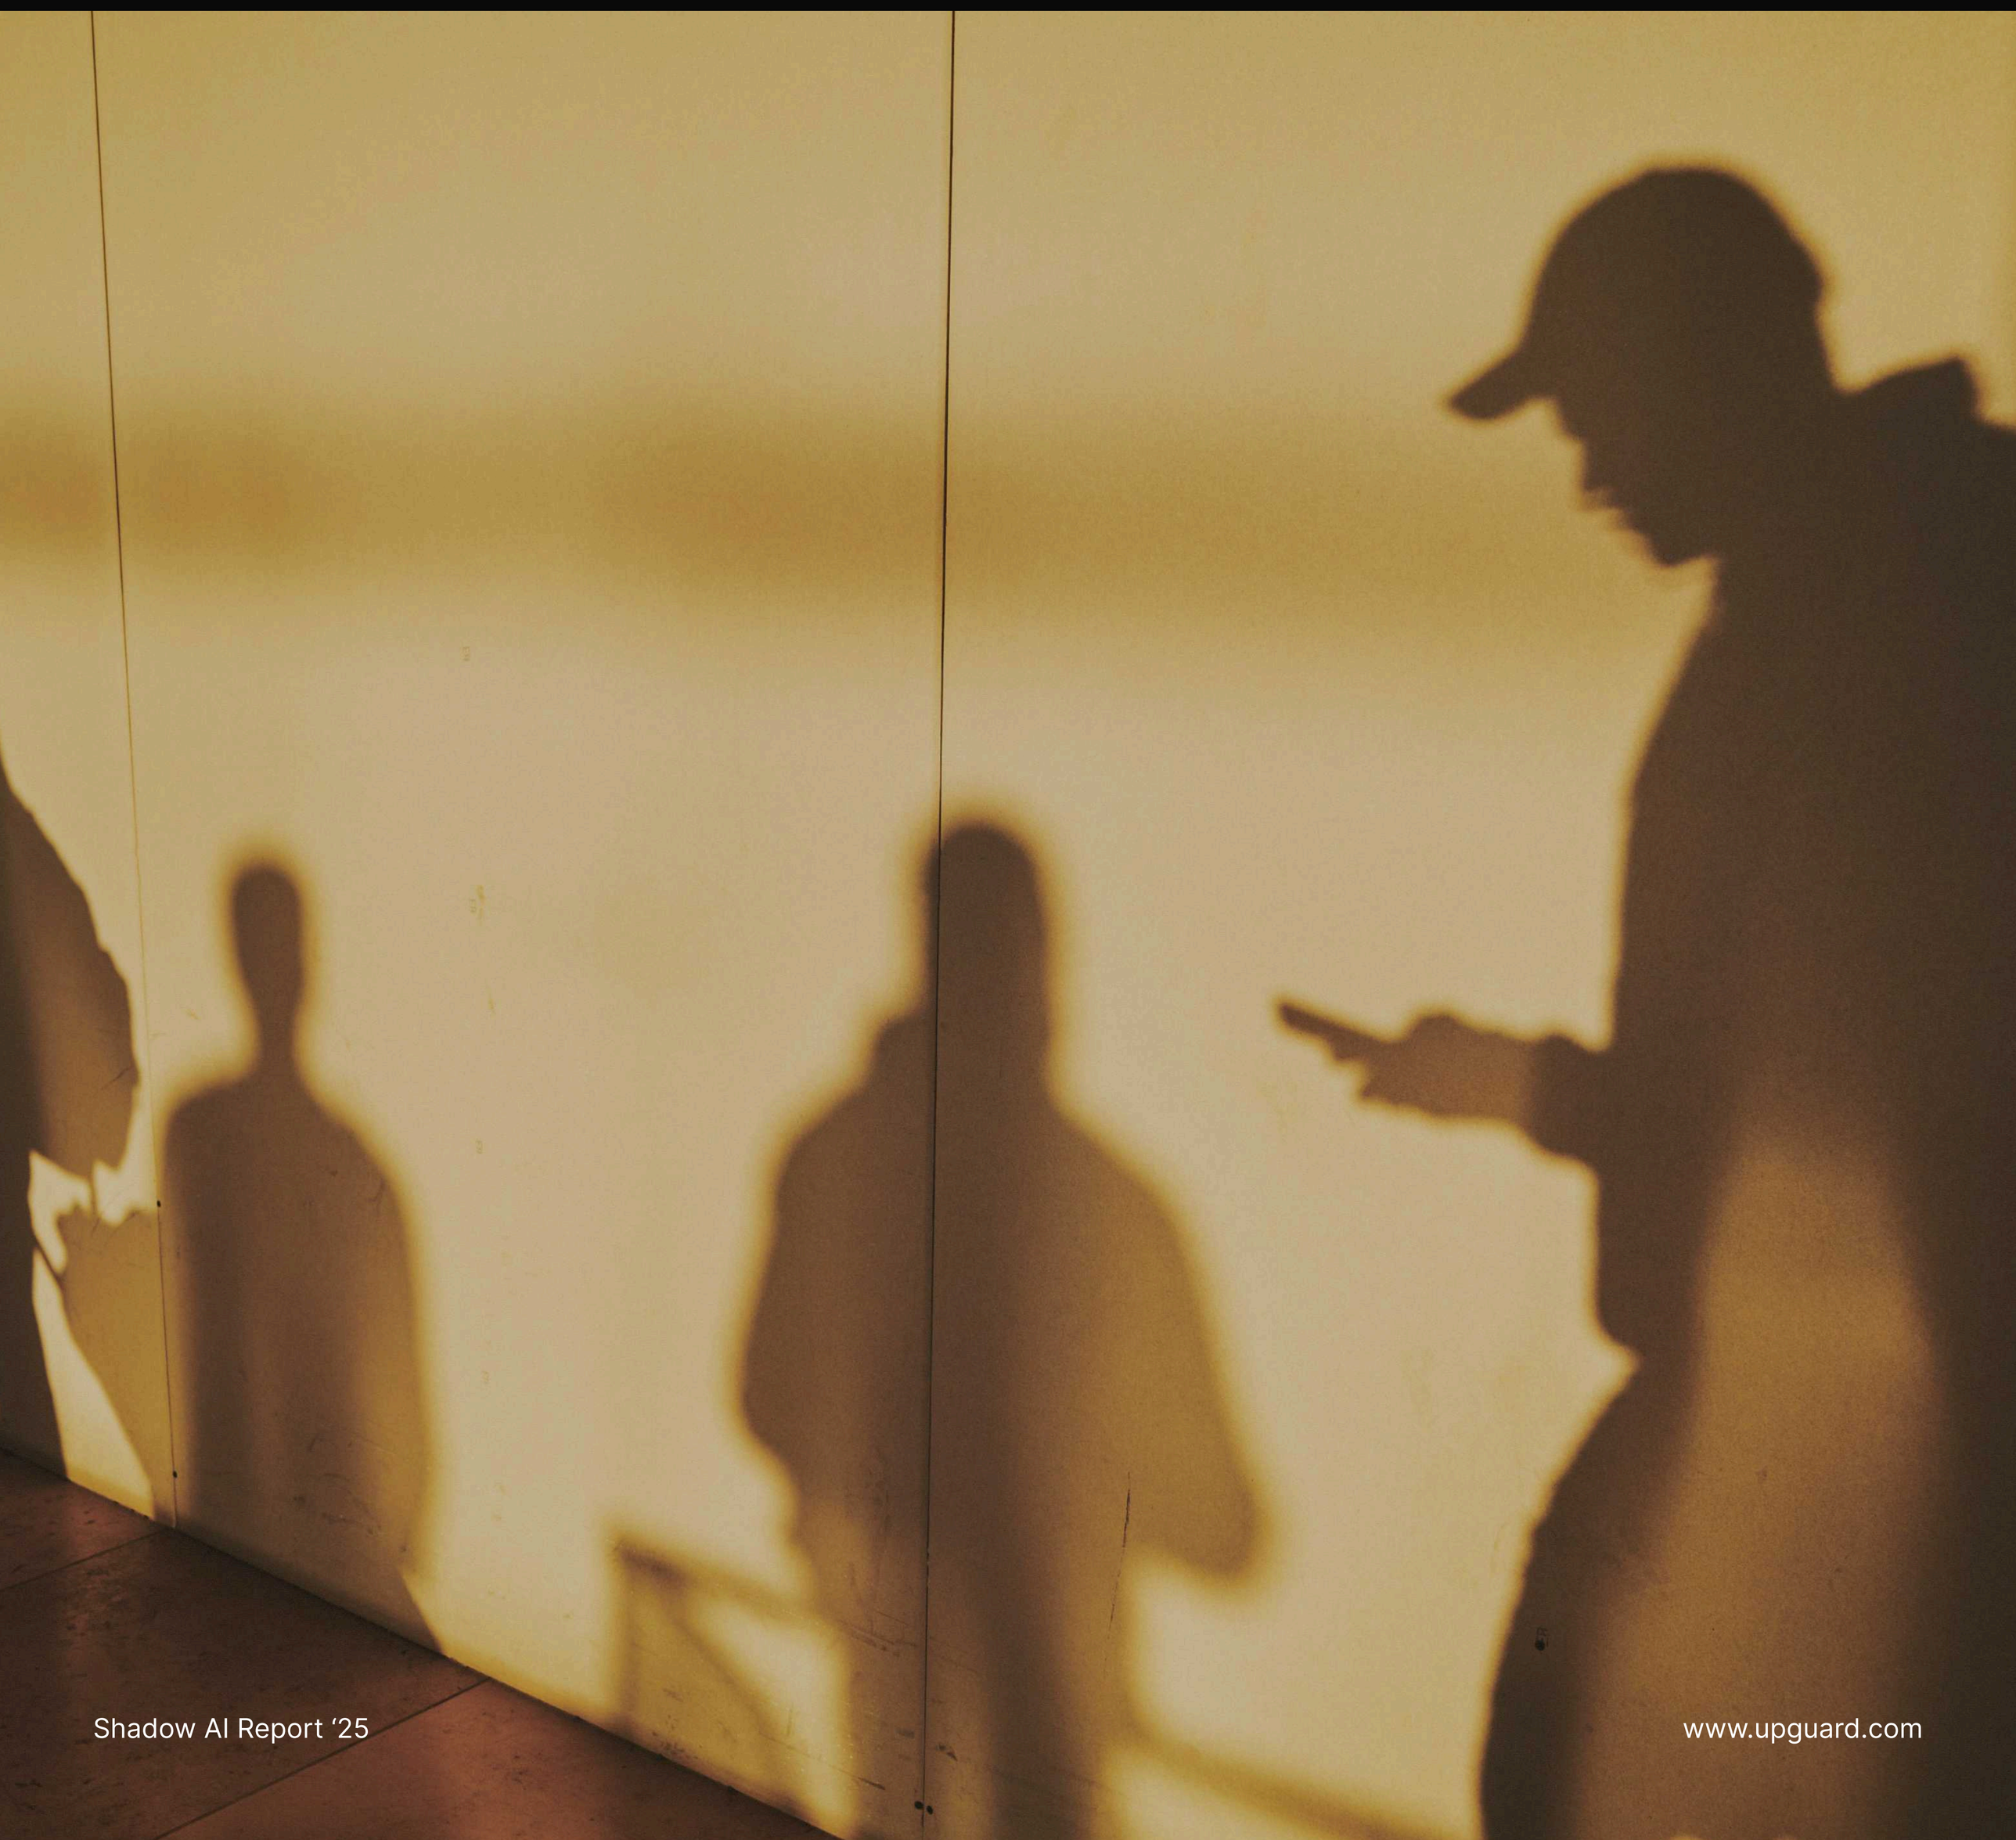

In the beginning there was ChatGPT, and millions of people used it, and it was good. Then they realized they could use it at work, and that was good too. Soon there were many AI tools, and they were good, and they were all being used to process corporate data without enterprise contracts and data privacy agreements. And that...that was not so good.

Previous studies have found 60-80% of employees use unapproved AI tools, also known as shadow AI. We wanted to go further and understand *why*.

Our first question was whether shadow AI correlates with other demographic and firmographic factors. Was shadow AI more prevalent in one industry or region? Are there particular departments within those companies creating more of a mess than others? Each of those factors do impact the overall rate of shadow AI, but the big picture is much simpler: shadow AI is a problem for every organization.

Our second question was whether shadow AI was the result of gaps in security awareness training. Here our results were more conclusive than expected, and counter-intuitive. People who reported better knowledge of AI security training also reported more regular use of shadow AI. Even more surprising, respondents to our survey of information security leaders reported even more shadow AI usage than others in the workforce.

Together, the effects of industry and expertise point to a theory of the AI power user: people who are active in the fast-moving conversation around AI and want to stay at the leading edge of tooling, even at the expense of corporate compliance.

Understanding shadow AI as an excess of enthusiasm should change our model for responding to it. Blocking tools does not work. In fact, almost half of the people we surveyed who encountered blocked tools found ways to continue accessing them.

Instead, the business needs to harness the curiosity driving shadow AI and put it to use. Just as shadow IT has been a consistent issue for decades, so has its solution: “CIOs need to make users feel comfortable about bringing their underground behavior into the light. Always try to help users figure out a safe and secure way to do whatever it is they’re trying to do,” wrote Ben Worthen in a 2007 article for CIO.com. Shadow AI is here to stay, but with the right tools for managing it, that just might be a good thing.

Methodology

To collect information about how people are using AI in the workplace and how information security leaders perceive employee use of AI, we conducted two surveys in August and September of 2024. All participants provided consent, were compensated a fair hourly wage for their time, and did not provide any personally identifying information to UpGuard.

Our survey of workers—defined as people who reported being employed—was administered through Prolific, an online platform used for previous studies of AI users like Microsoft's "The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers." Employees were screened to require non-null values for the industry and size of their employer.

Our survey of information security leaders was performed by Dynata, "the world's largest and highest-quality first-party data company." Participants were screened as "security leaders" based on self-reported job title and through questions designed to identify users in strategic roles.

Through Prolific we surveyed 500 participants in the US and 500 in the UK. Through Dynata we surveyed 500 security leaders in the US, Canada, APAC (Australia, New Zealand, and Singapore), and India.

Everyone’s using AI, whether it’s allowed or not.

Across companies of all sizes, geographies, and industries, employees are regularly using shadow AI, including information security leaders.

In our survey of 1,500 workers around the world, we found that shadow AI is everywhere—with some interesting wrinkles. Regardless of company size, geography, industry, employee function or seniority, a sizable majority of workers use AI tools at work that they know are not approved.

Within that overall trend, however, we also found that some factors correlate with higher rates of shadow AI. At the company level, those in regions and industries with a greater emphasis on technological innovation also see more usage of unapproved AI tools. Within companies, the likelihood increases depending on department and employee seniority.

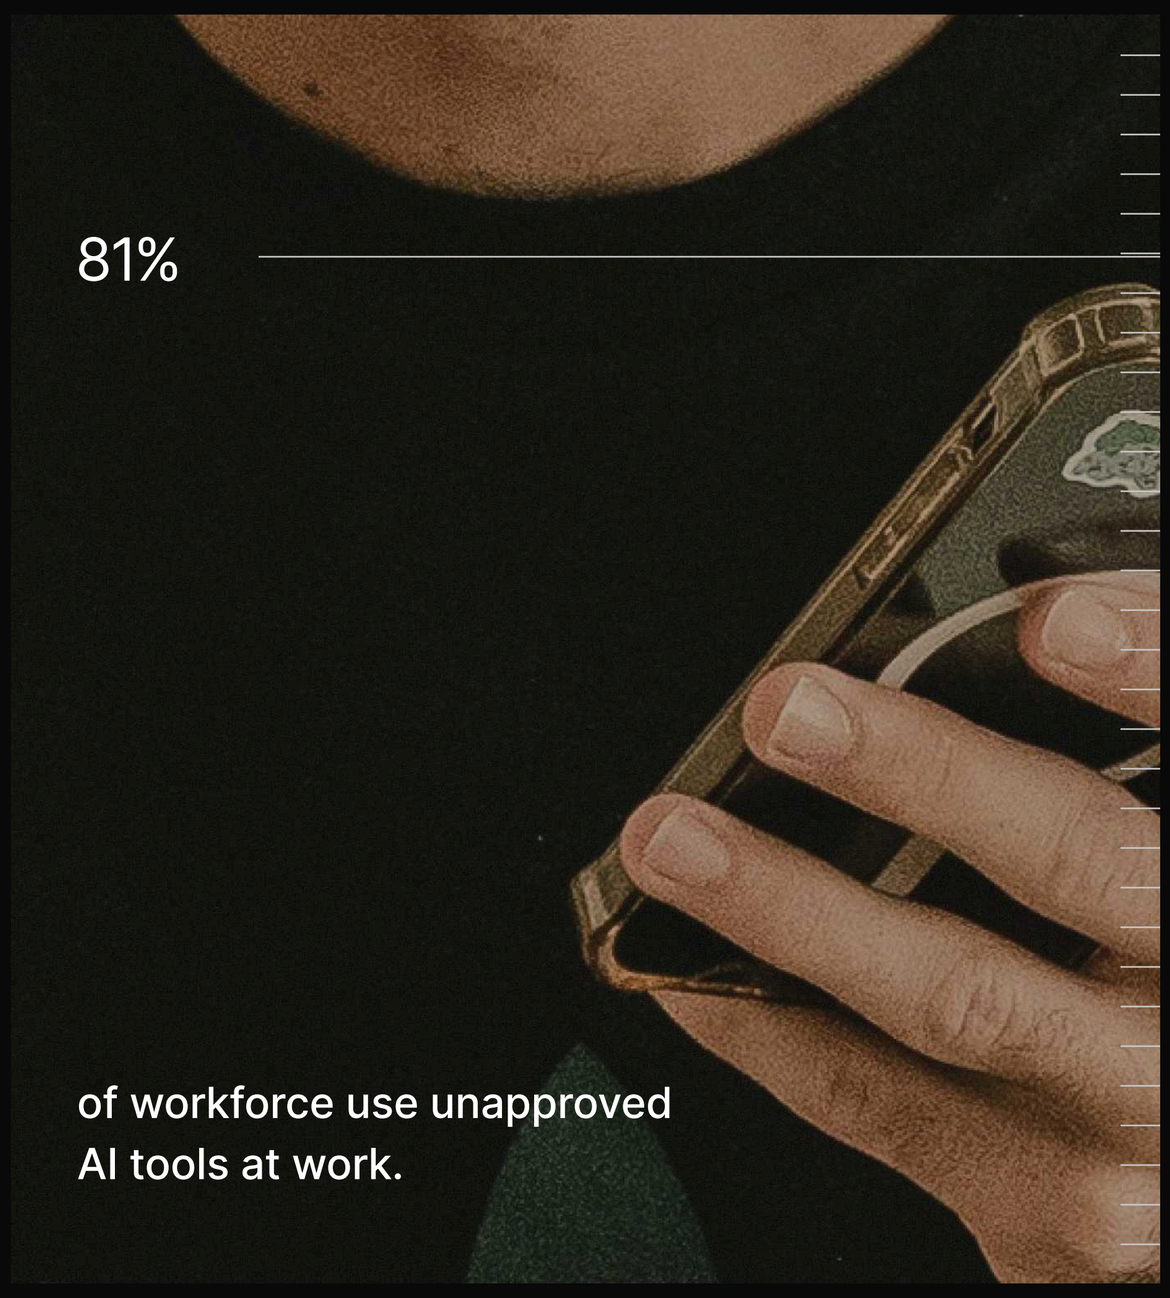

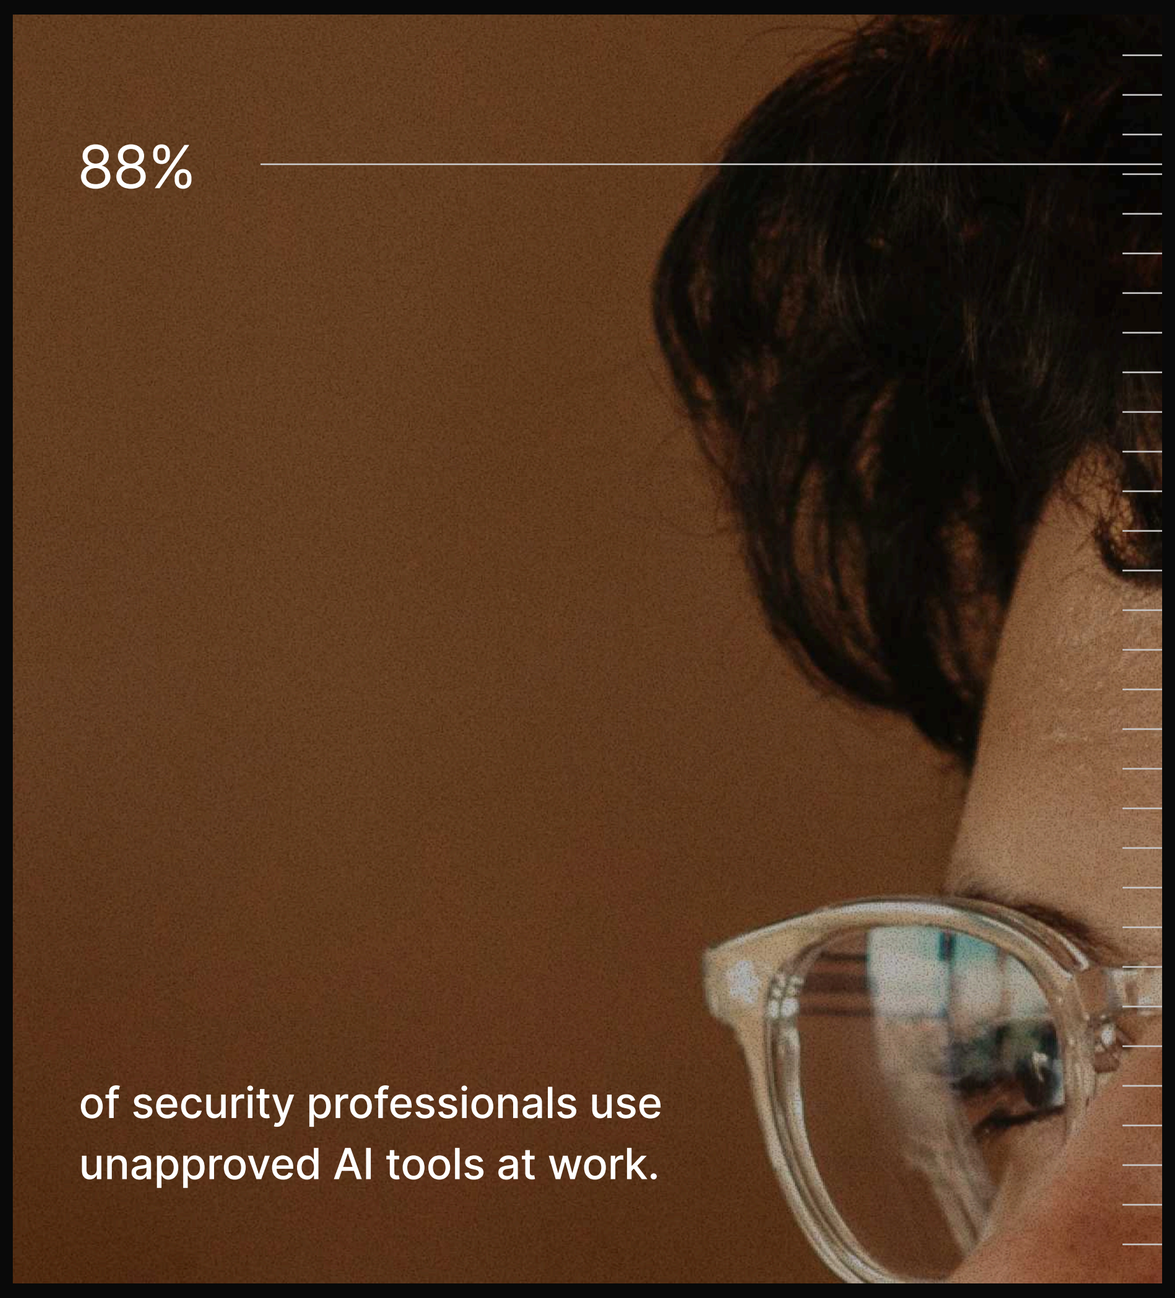

Workforce Usage of Shadow AI

Across the workforce, shadow AI use is widespread. Half of all workers said they use unapproved AI tools as part of their regular workflow, another third turn to them occasionally, and only 17% rely entirely on company-provided options.

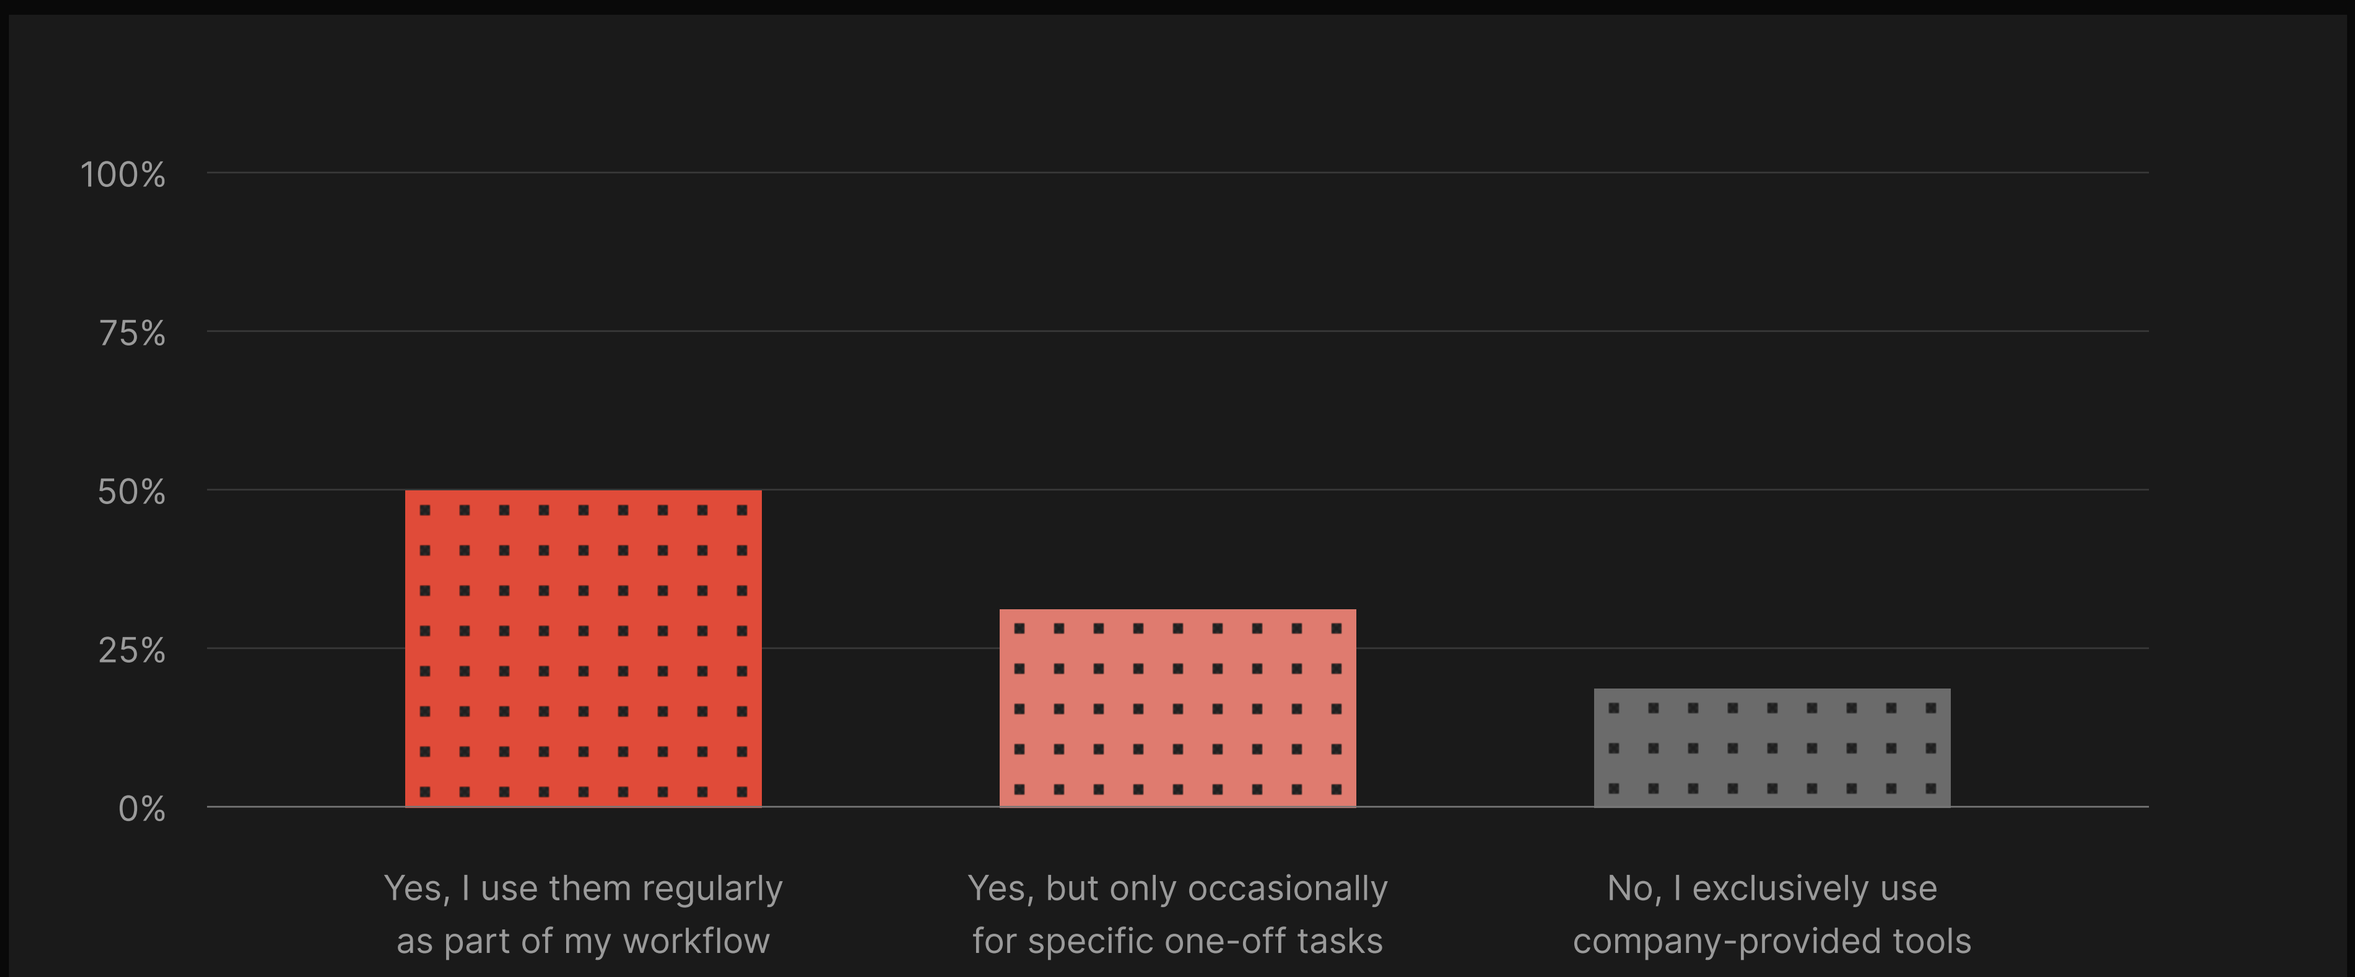

Security Leaders Usage of Shadow AI

Information security leaders were more likely to report that they use unapproved AI tools, and far more likely to report doing so regularly. 88% of security leaders reported using shadow AI, but were 33% more likely to do so regularly than other workers.

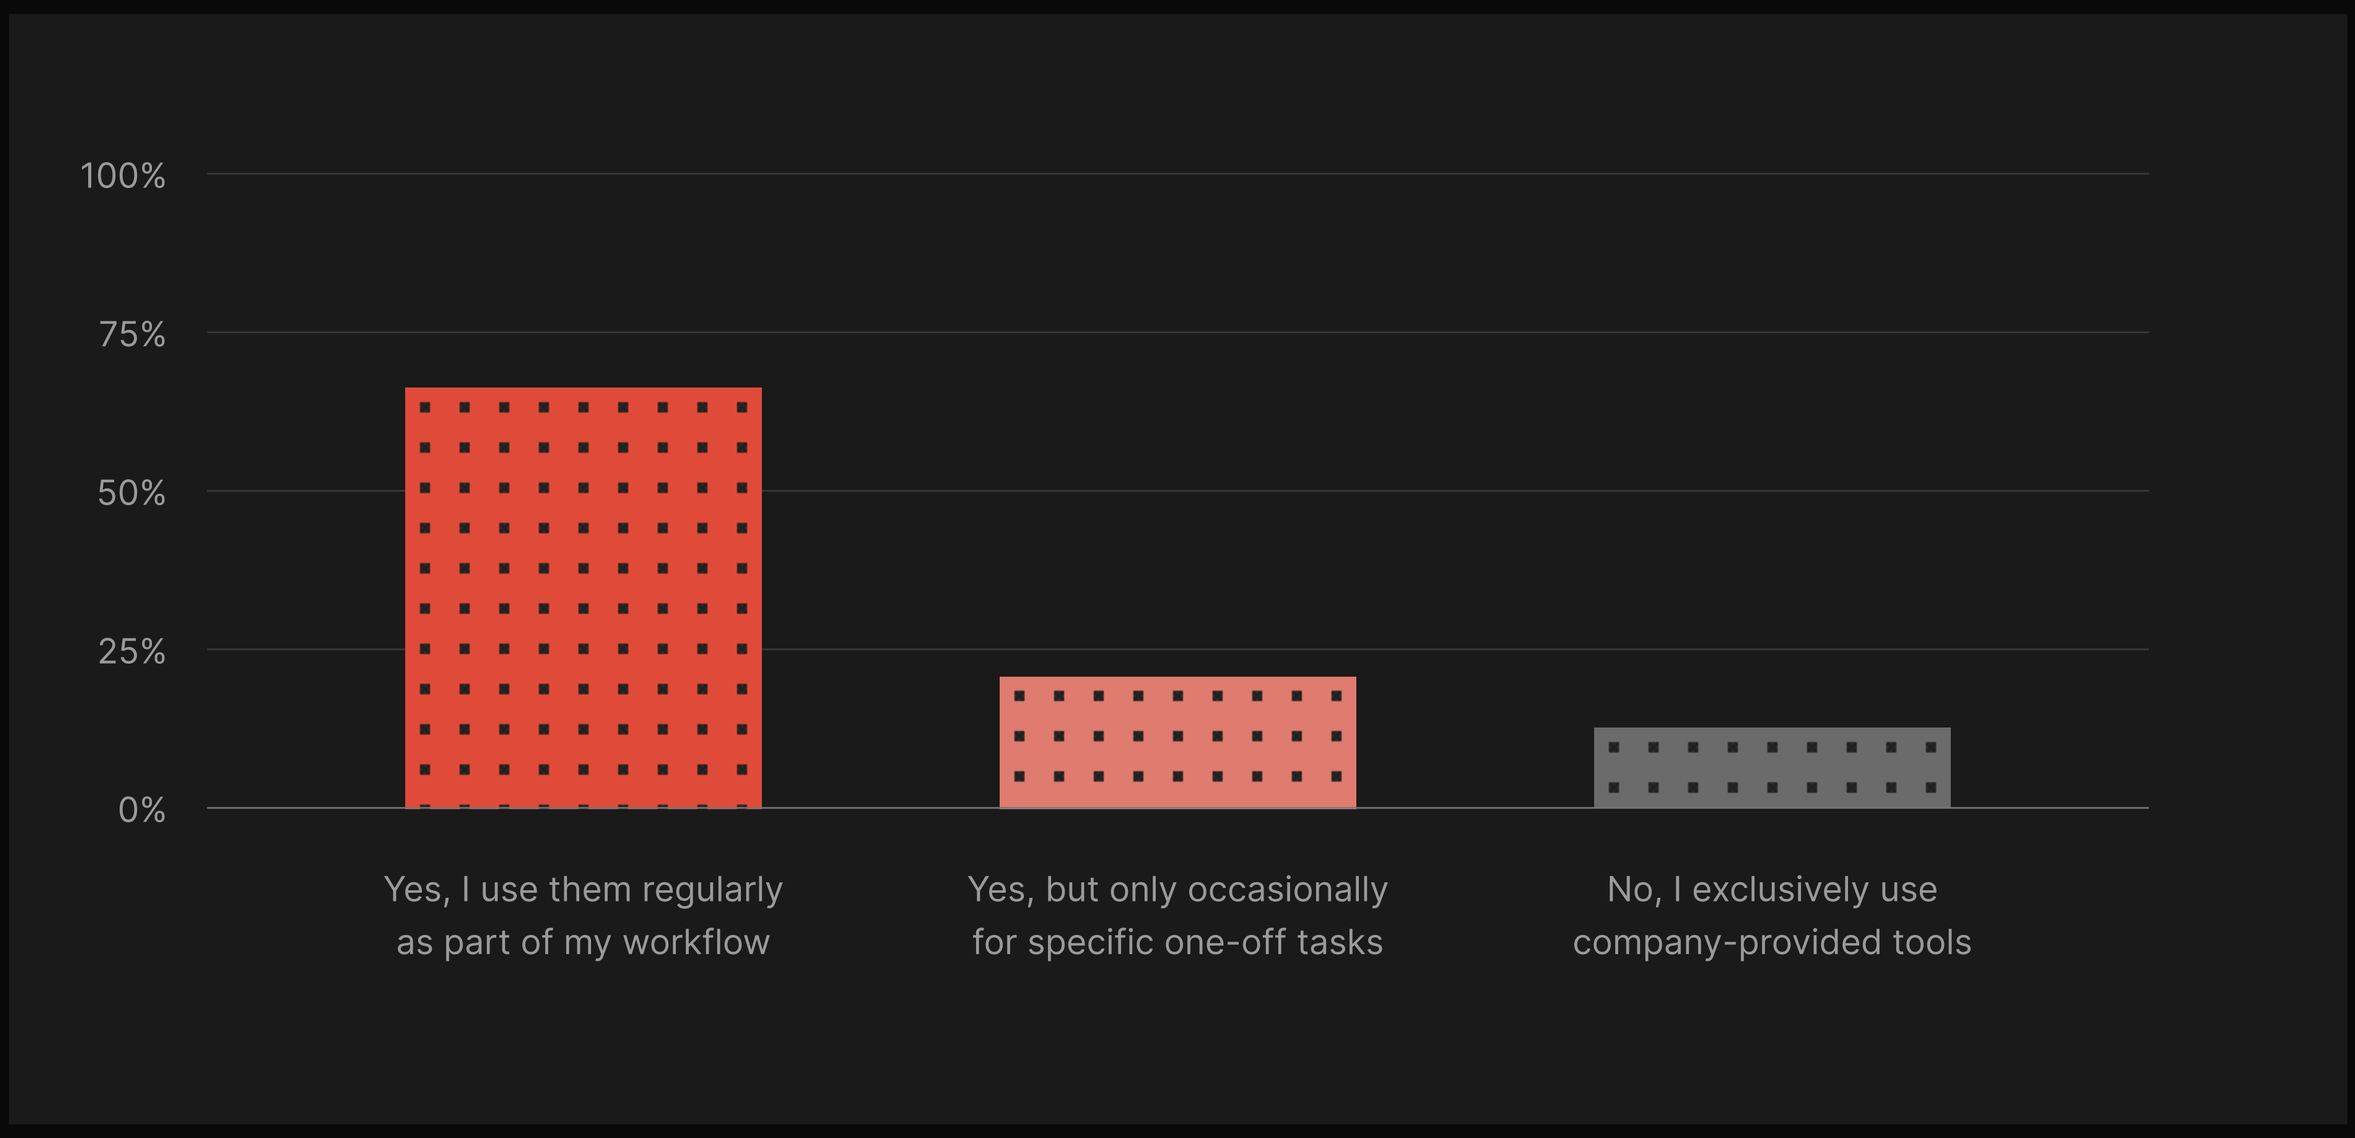

Security Leaders Usage of Shadow AI by Region

Comparing the results from security leaders sliced by region, some groups are even more likely to be habitual users of shadow AI. The overall mean of 88% using shadow AI climbs to 93% in the US with an equivalent increase in the number using it frequently.

In APAC and Canada, where AI regulation laws are in progress, security leaders report lower rates of shadow AI than in the “pro-innovation” regions of the US and India. These results suggest that cultural norms and government regulation may have some effect on individuals’ willingness to incur the risk of shadow AI.

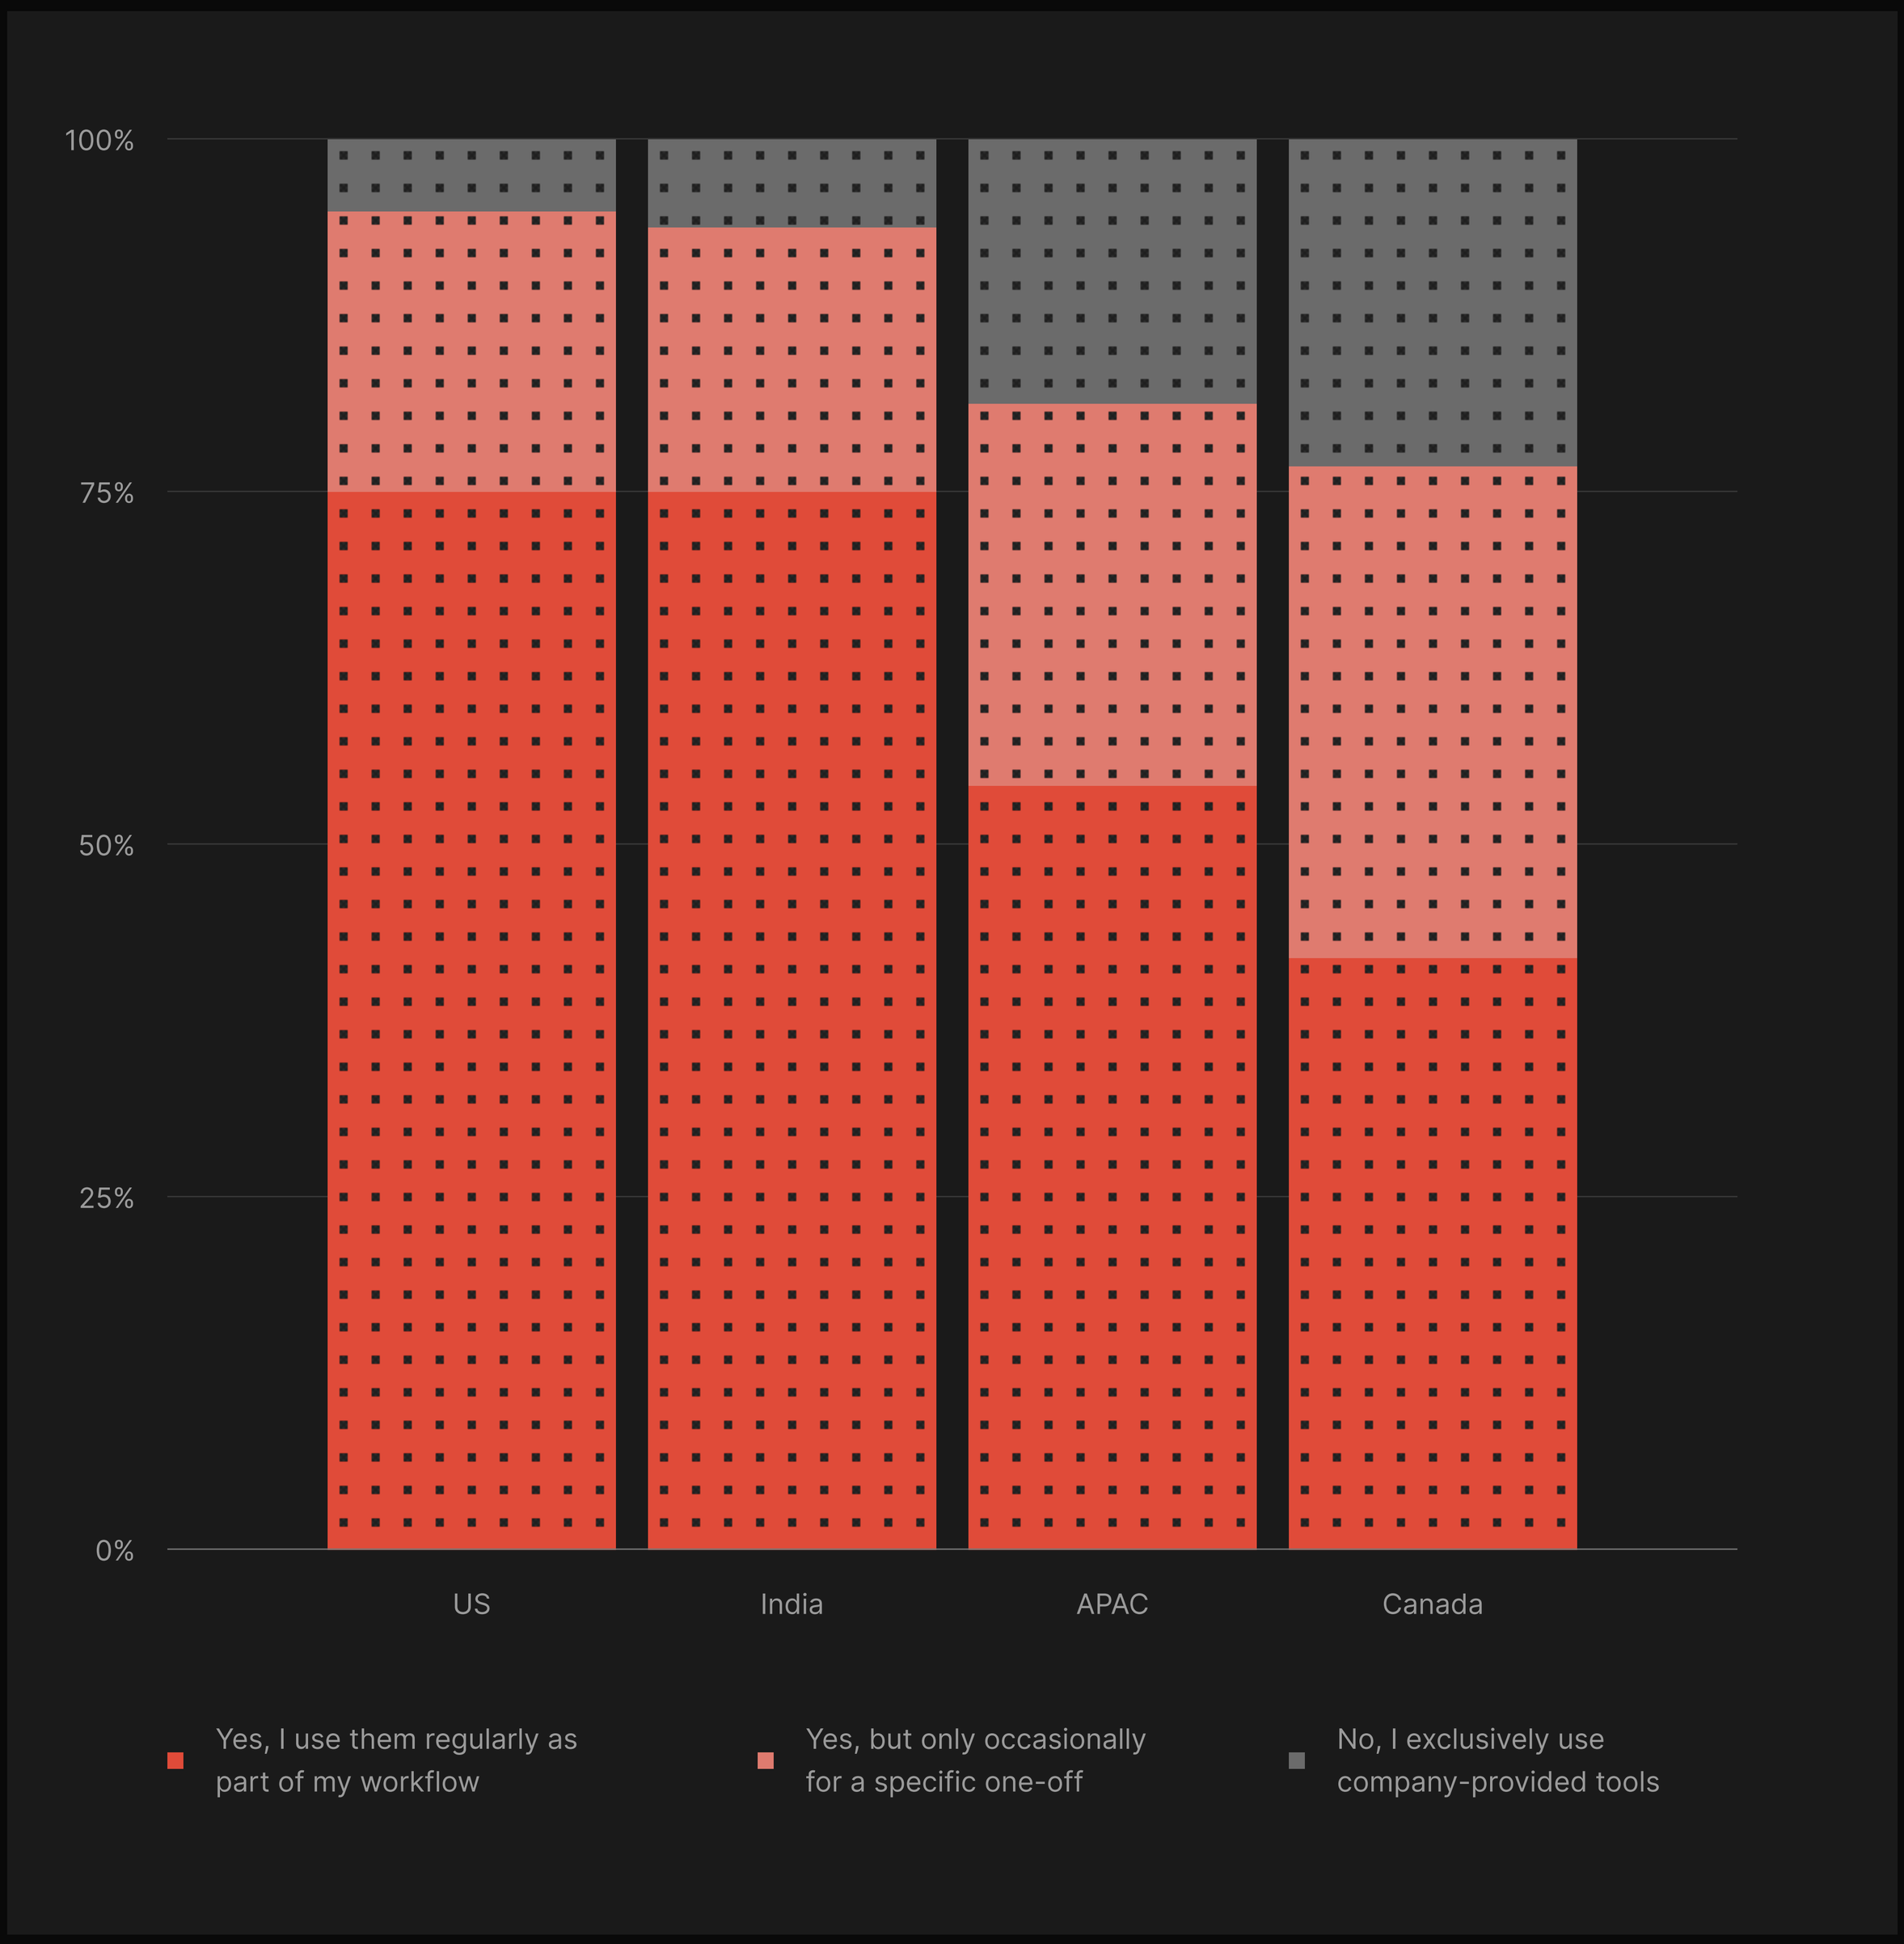

Usage of Shadow AI by Company Size

Workers at the very largest companies reported slightly lower rates of shadow AI, but across other company sizes the differences were within the margin of error.

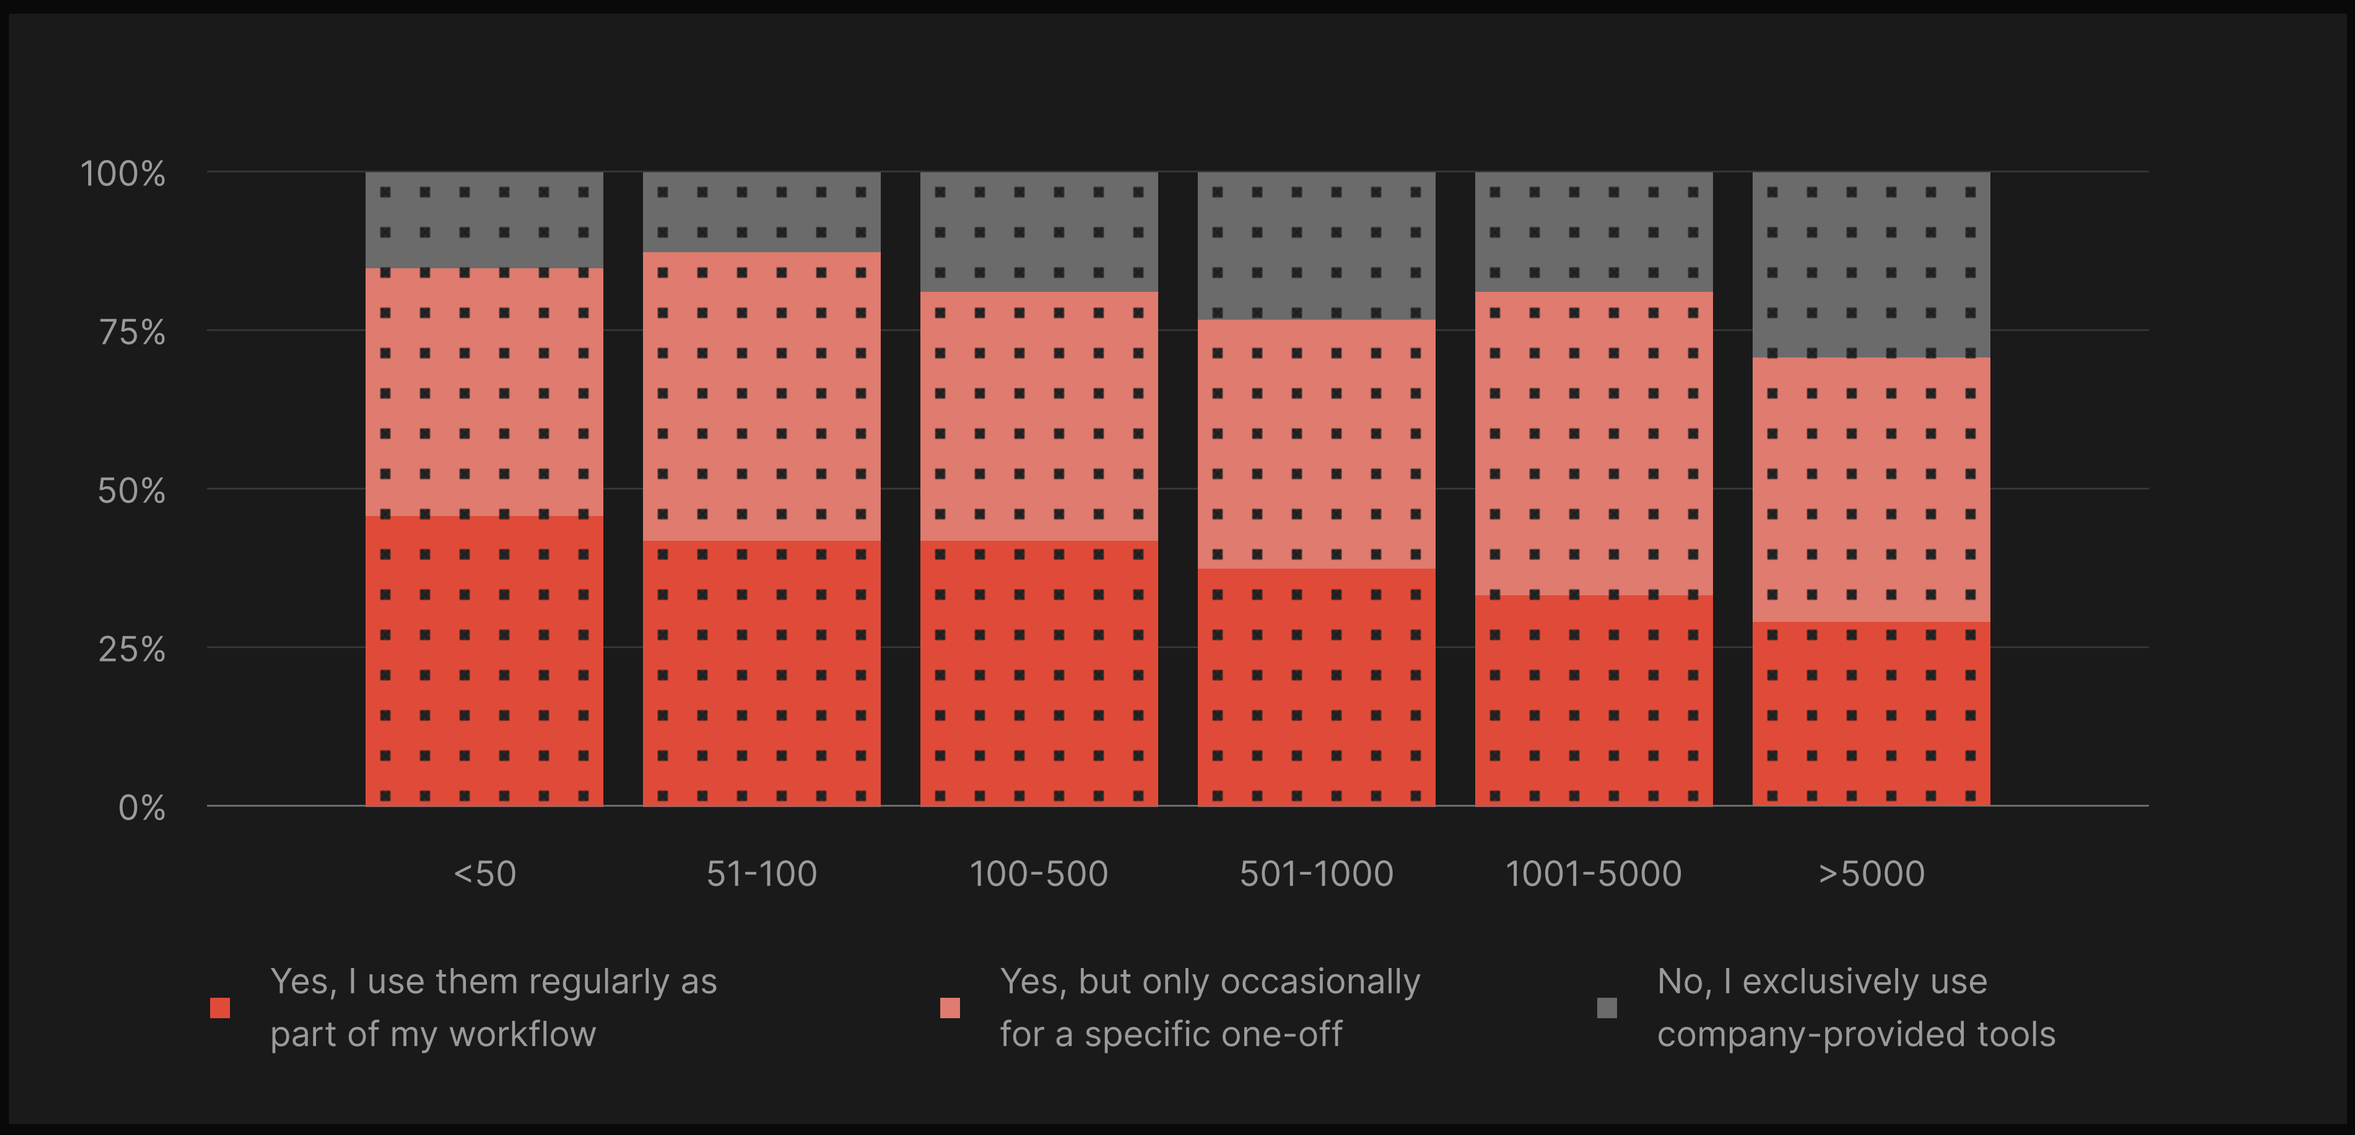

Usage of Shadow AI by Industry

Industry has mild effects on shadow AI usage, with some larger effects on the frequency of regularly using such tools. 10% more Information and Finance employees, for example, regularly use shadow AI than those in several other industries.

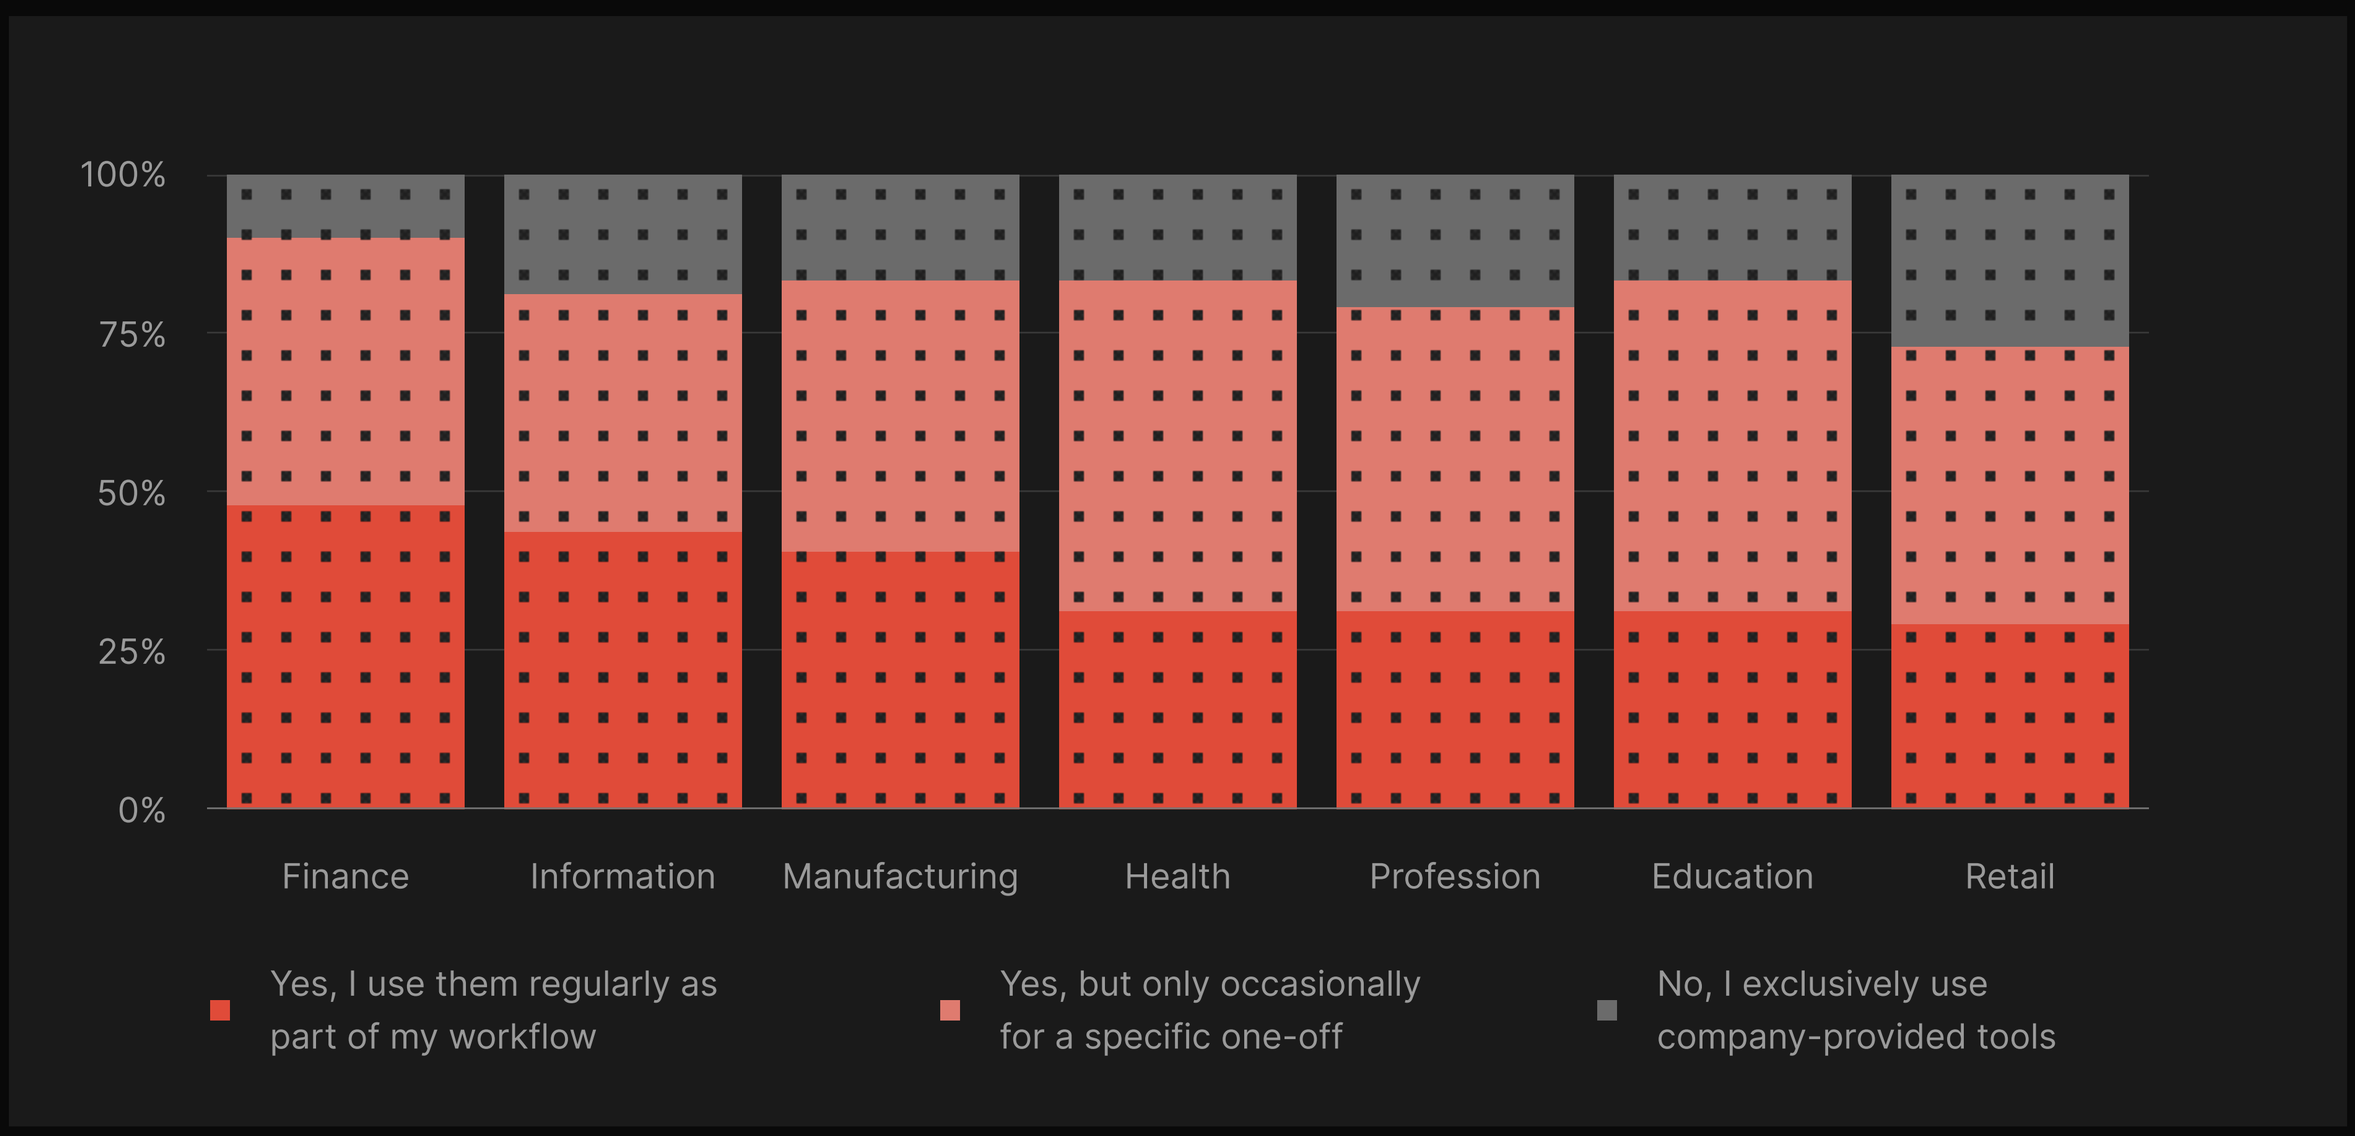

Usage of Shadow AI by Role

At every level of seniority, we find substantial use of shadow AI. Executives reported by far the highest rate of regular use of shadow AI. Whether that is surprising depends on one's opinions about executives.

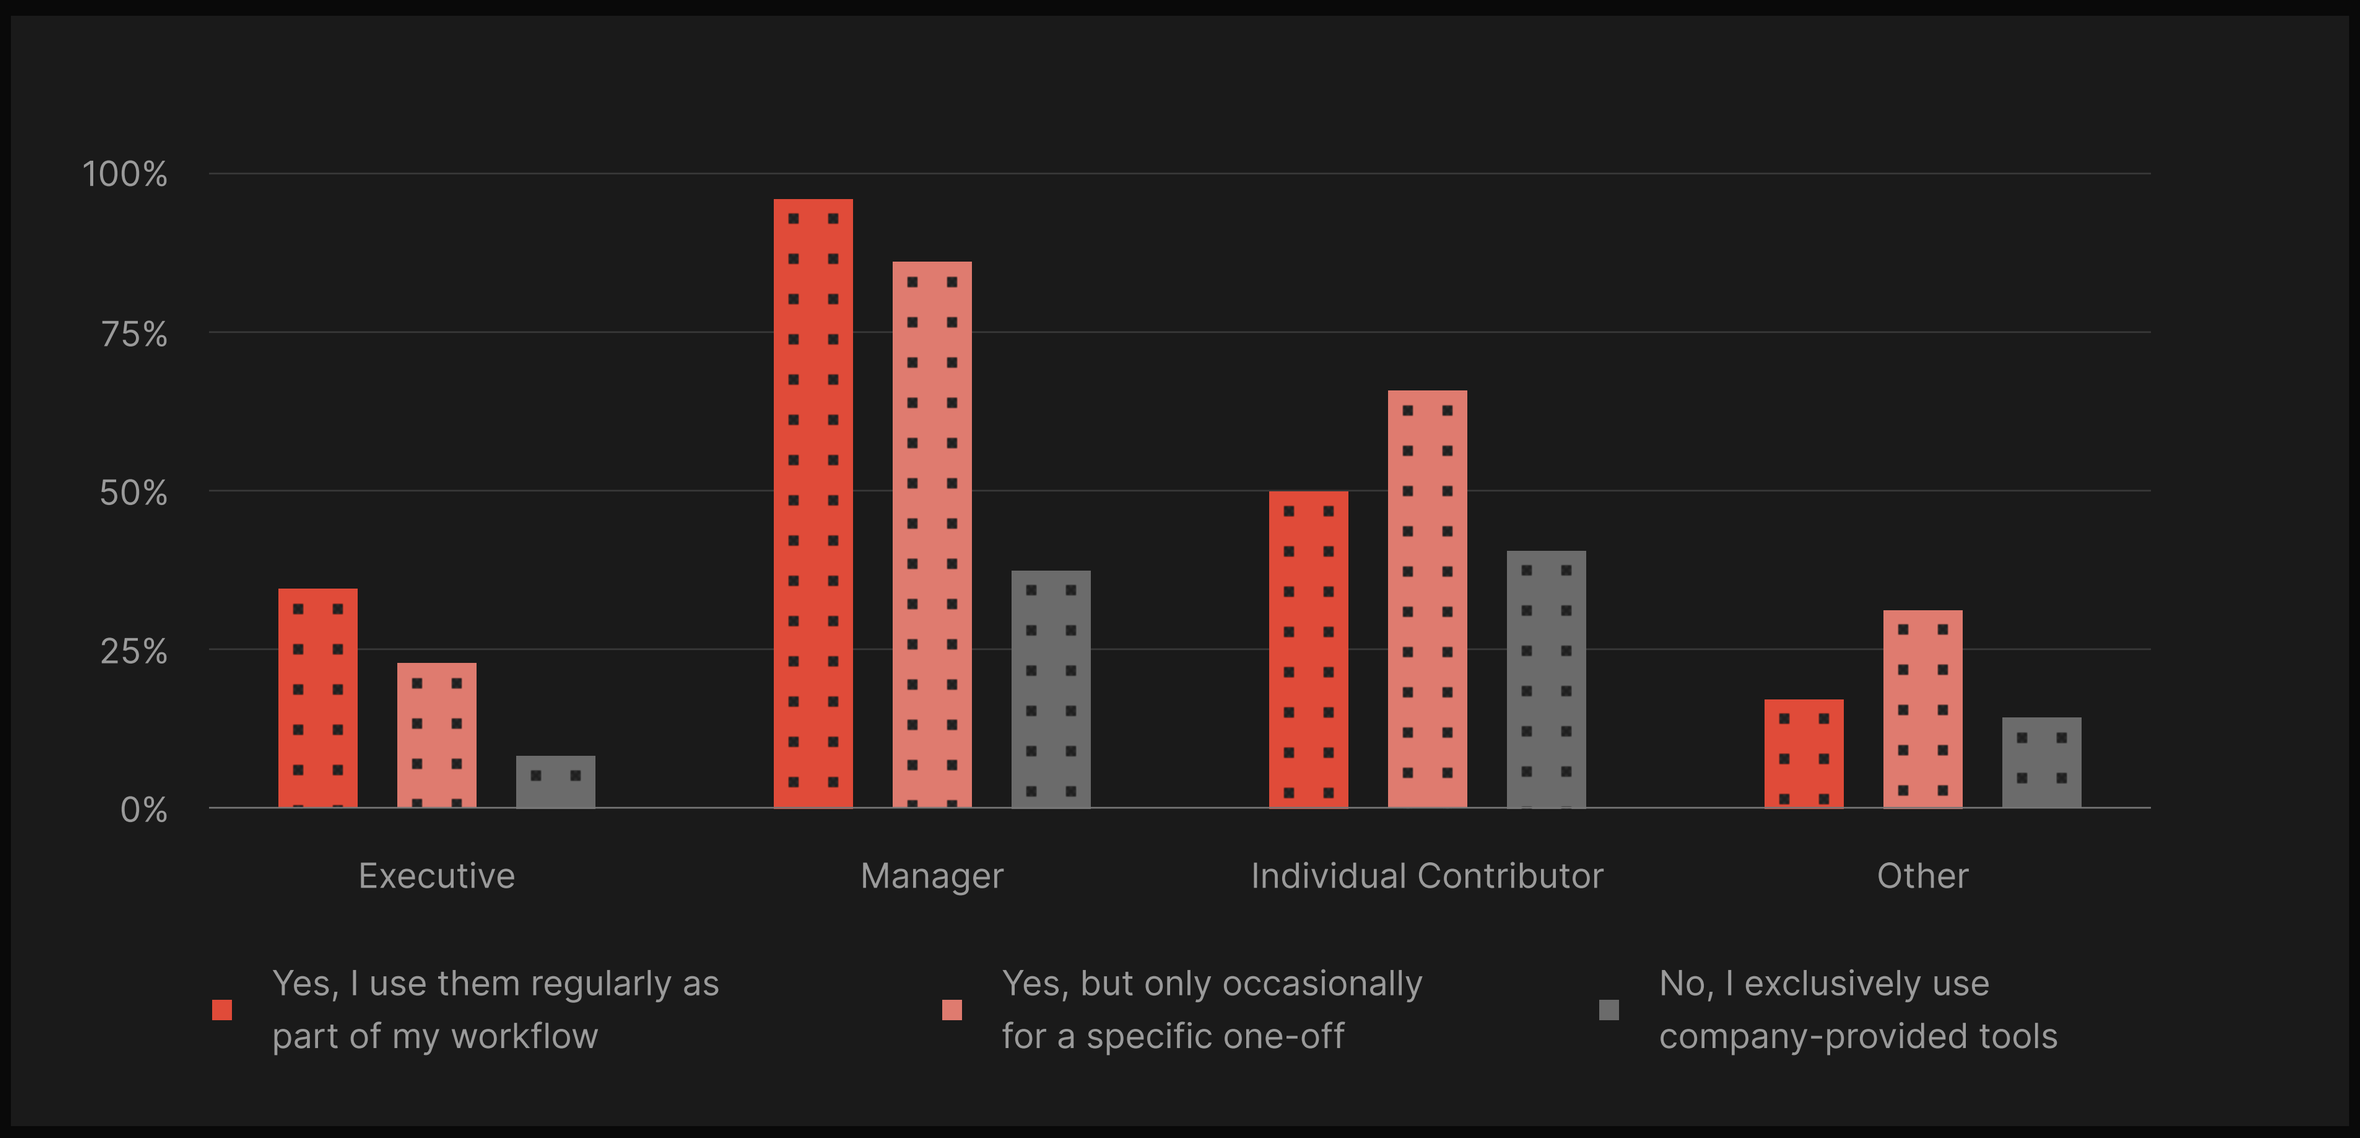

Usage of Shadow AI by Department

Sales and Marketing are commonly perceived as the most likely to fail phishing simulations, but many departments, including HR, show comparable levels of shadow AI. Only Ops, Finance, and Support have significantly less of risk from unapproved AI tools.

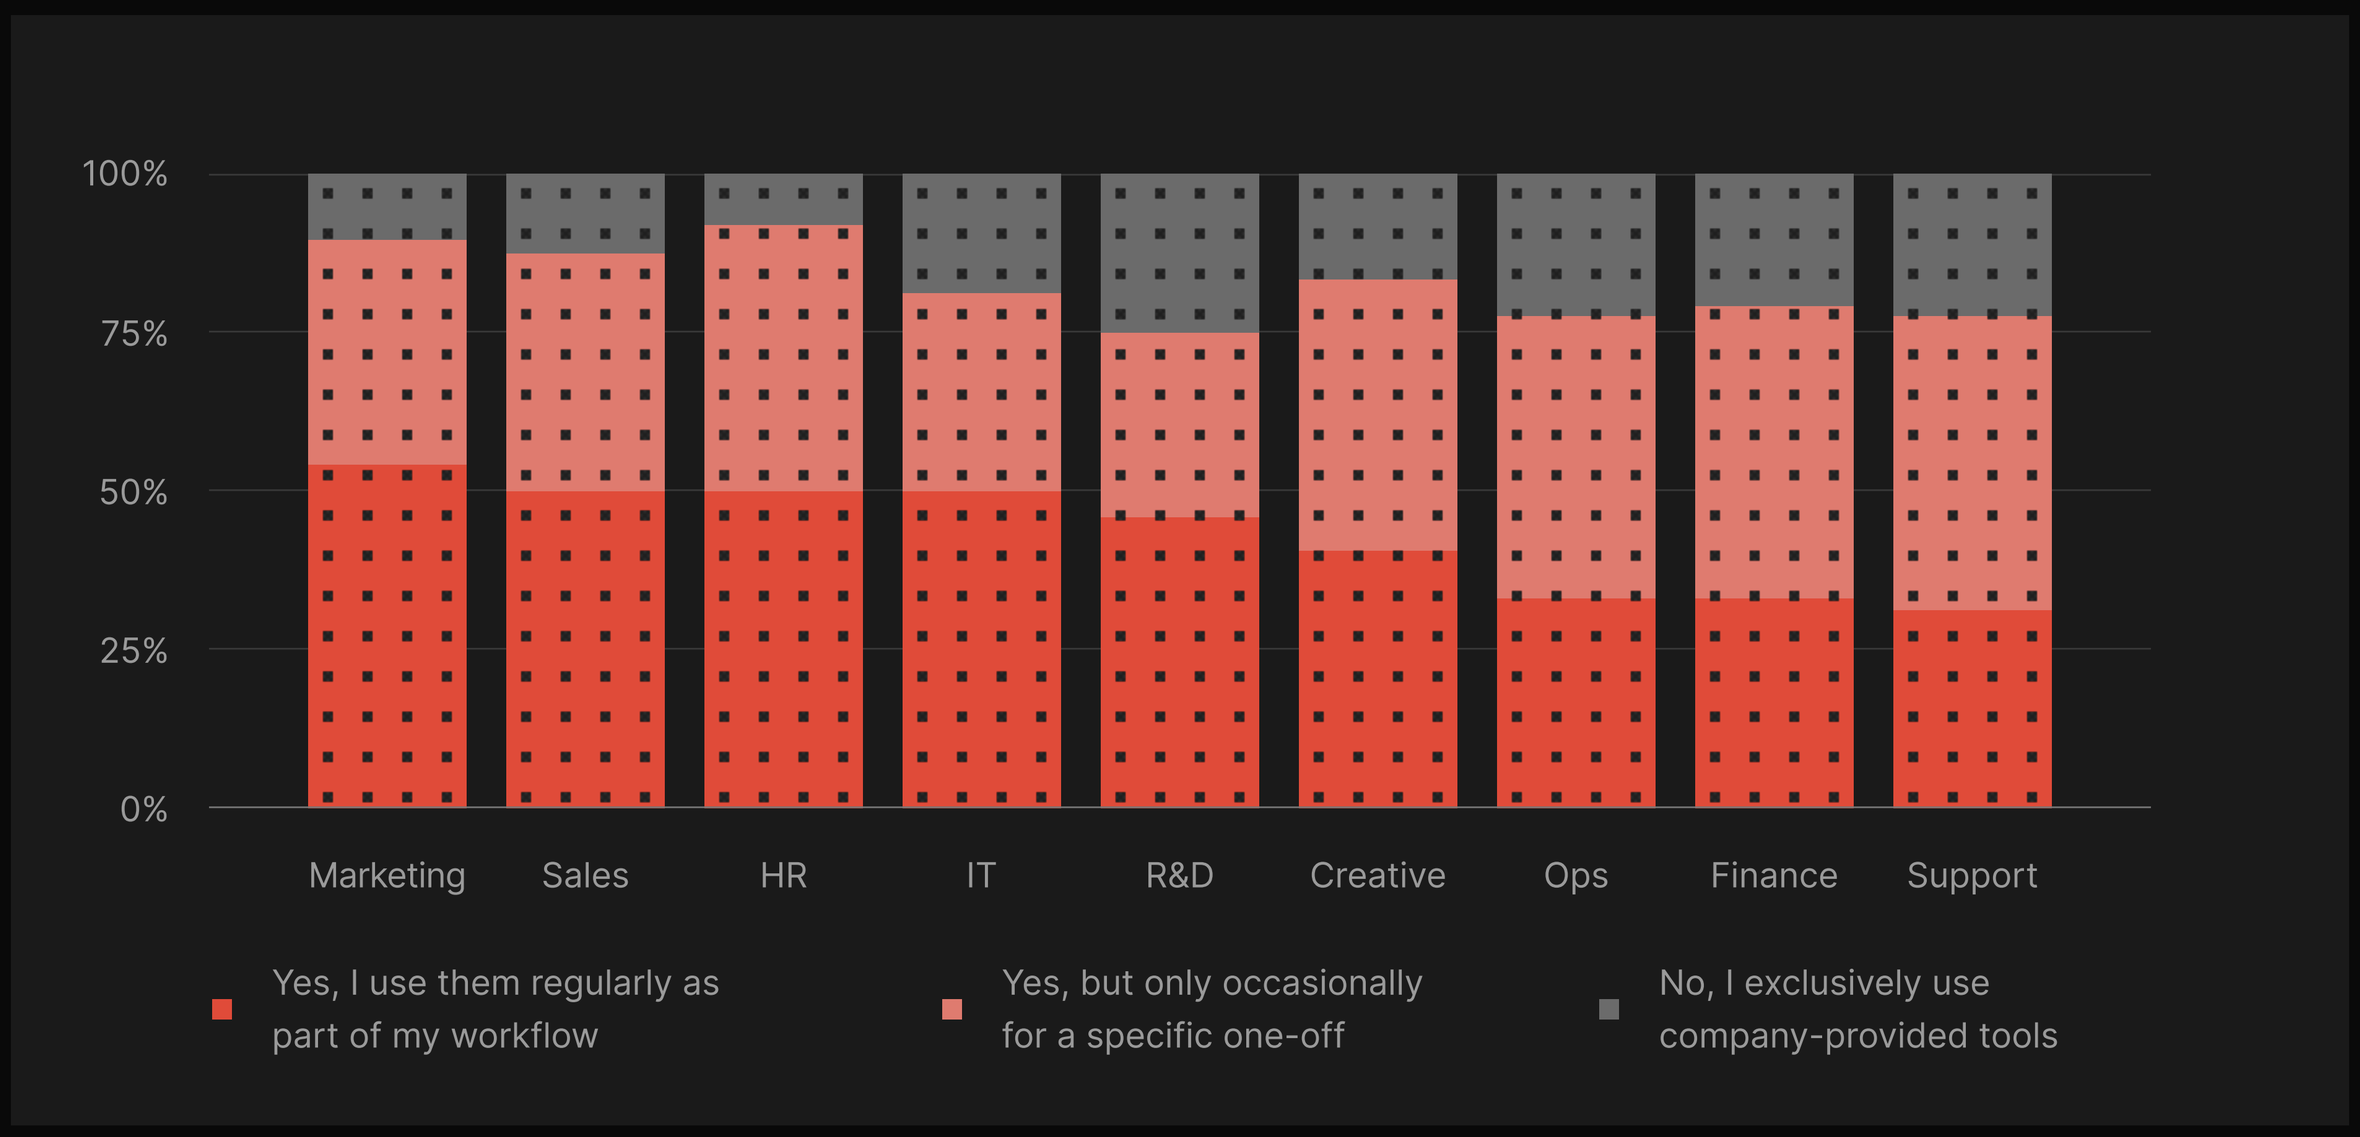

The harder you try to control it, the worse it gets.

The more employees understand the risks of AI, the more likely they are to use unapproved AI tools.

Like previous studies of shadow AI, our survey found that a large percentage of employees are using unapproved AI tools at work. Reproducing that finding enabled us to go beyond previous studies by asking other questions about those users' knowledge, beliefs, and behaviors related to AI.

Our initial hypothesis was that employees used shadow AI because of a knowledge gap. AI is relatively new, and employees might be using unapproved AI tools because they are unaware of the extent to which that is perceived as a risk.

To test this hypothesis, we asked 1,000 employees three questions related to AI governance and security awareness training: 'Could they recall their company's AI usage policy', 'Could they recall employer provided AI safety training', and 'Did they know how their AI tools used inputs to train models.'

Surprisingly, across all of these measures of AI literacy, we found a positive correlation between users reporting that they understood AI security requirements and that they regularly used unapproved AI tools.

This data suggests that as employees' knowledge of AI risks increases, so does their confidence in making judgments about that risk—even at the expense of following company policies. While this does not mean we should avoid AI security awareness training, it certainly indicates it is not sufficient, and that such programs need new approaches in order to succeed.

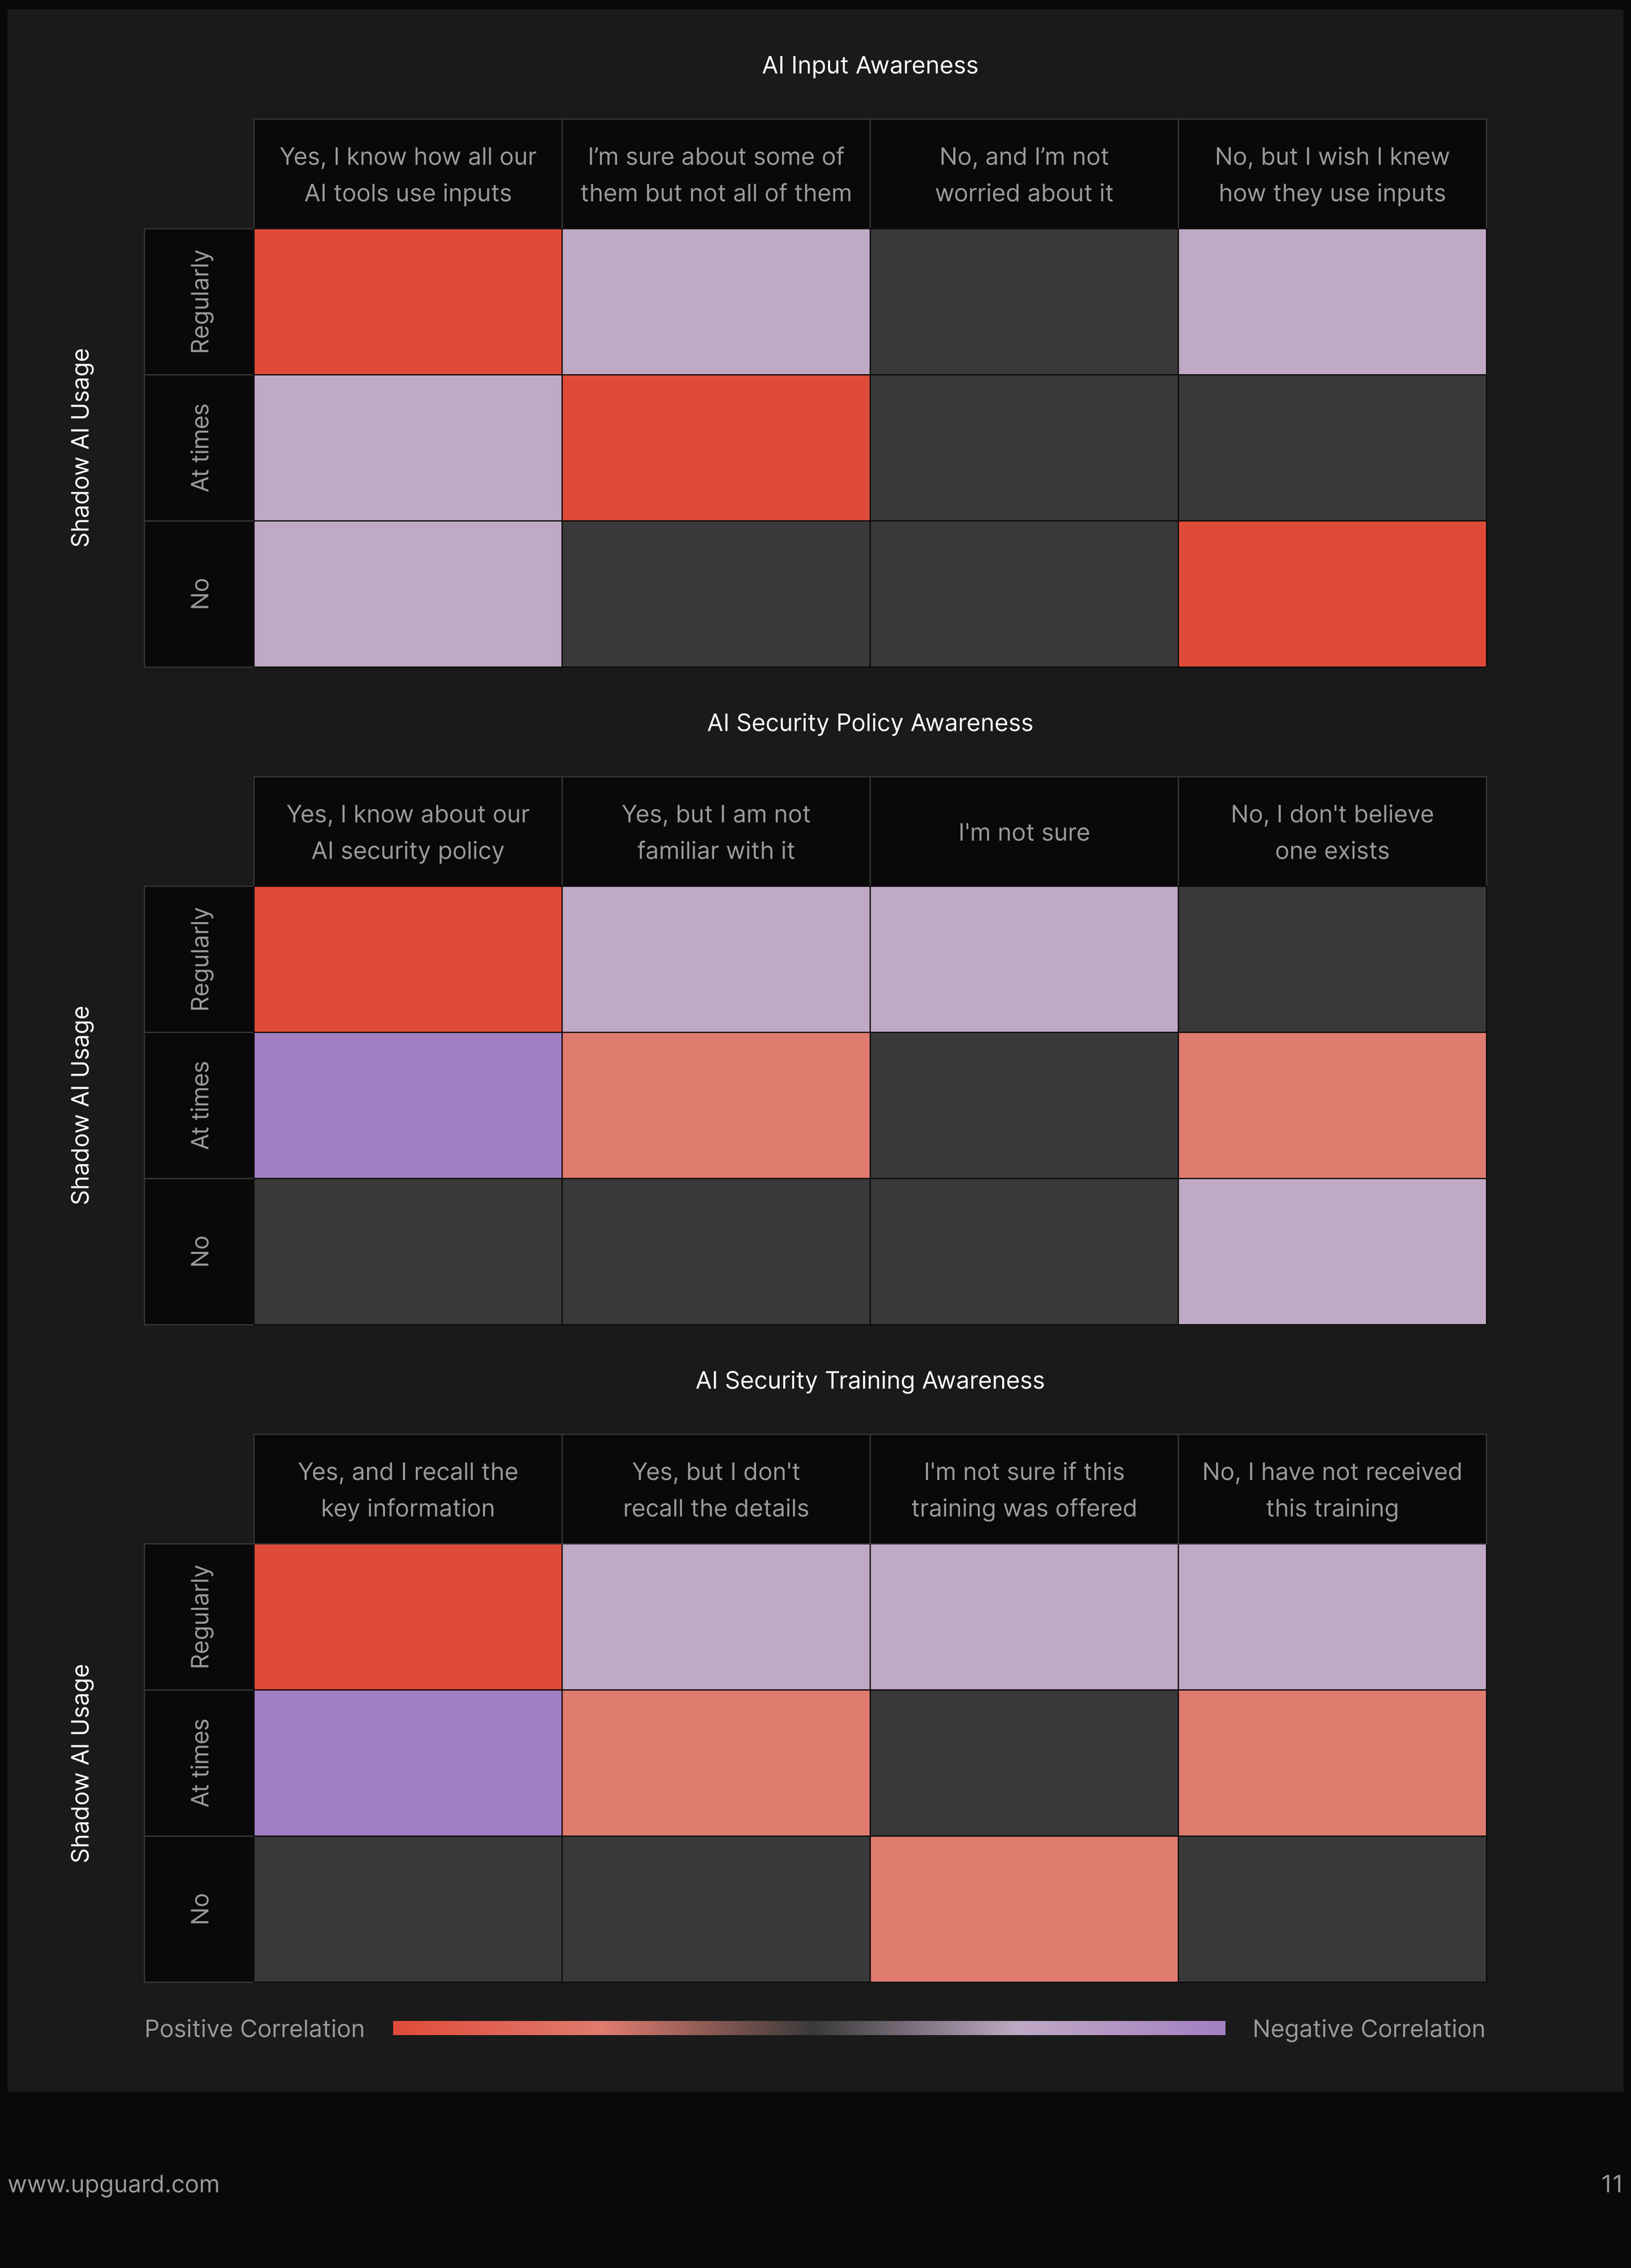
AI Input Awareness
Yes, I know how all our AI tools use inputs
I'm sure about some of them but not all of them
No, and I'm not worried about it
No, but I wish I knew how they use inputs
Shadow AI Usage
Regularly
At times
No
AI Security Policy Awareness
Yes, I know about our AI security policy
Yes, but I am not familiar with it
I'm not sure
No, I don't believe one exists
Shadow AI Usage
Regularly
At times
No
AI Security Training Awareness
Yes, and I recall the key information
Yes, but I don't recall the details
I'm not sure if this training was offered
No, I have not received this training
Shadow AI Usage
Regularly
At times
No
Positive Correlation
Negative Correlation

Shadow AI Use by Level of AI Awareness

Combining the results from the three questions about workers' knowledge of AI governance, we found a general correlation between increasing knowledge and increasing shadow AI. This positive correlation was both the strongest in effect size and statistical significance.

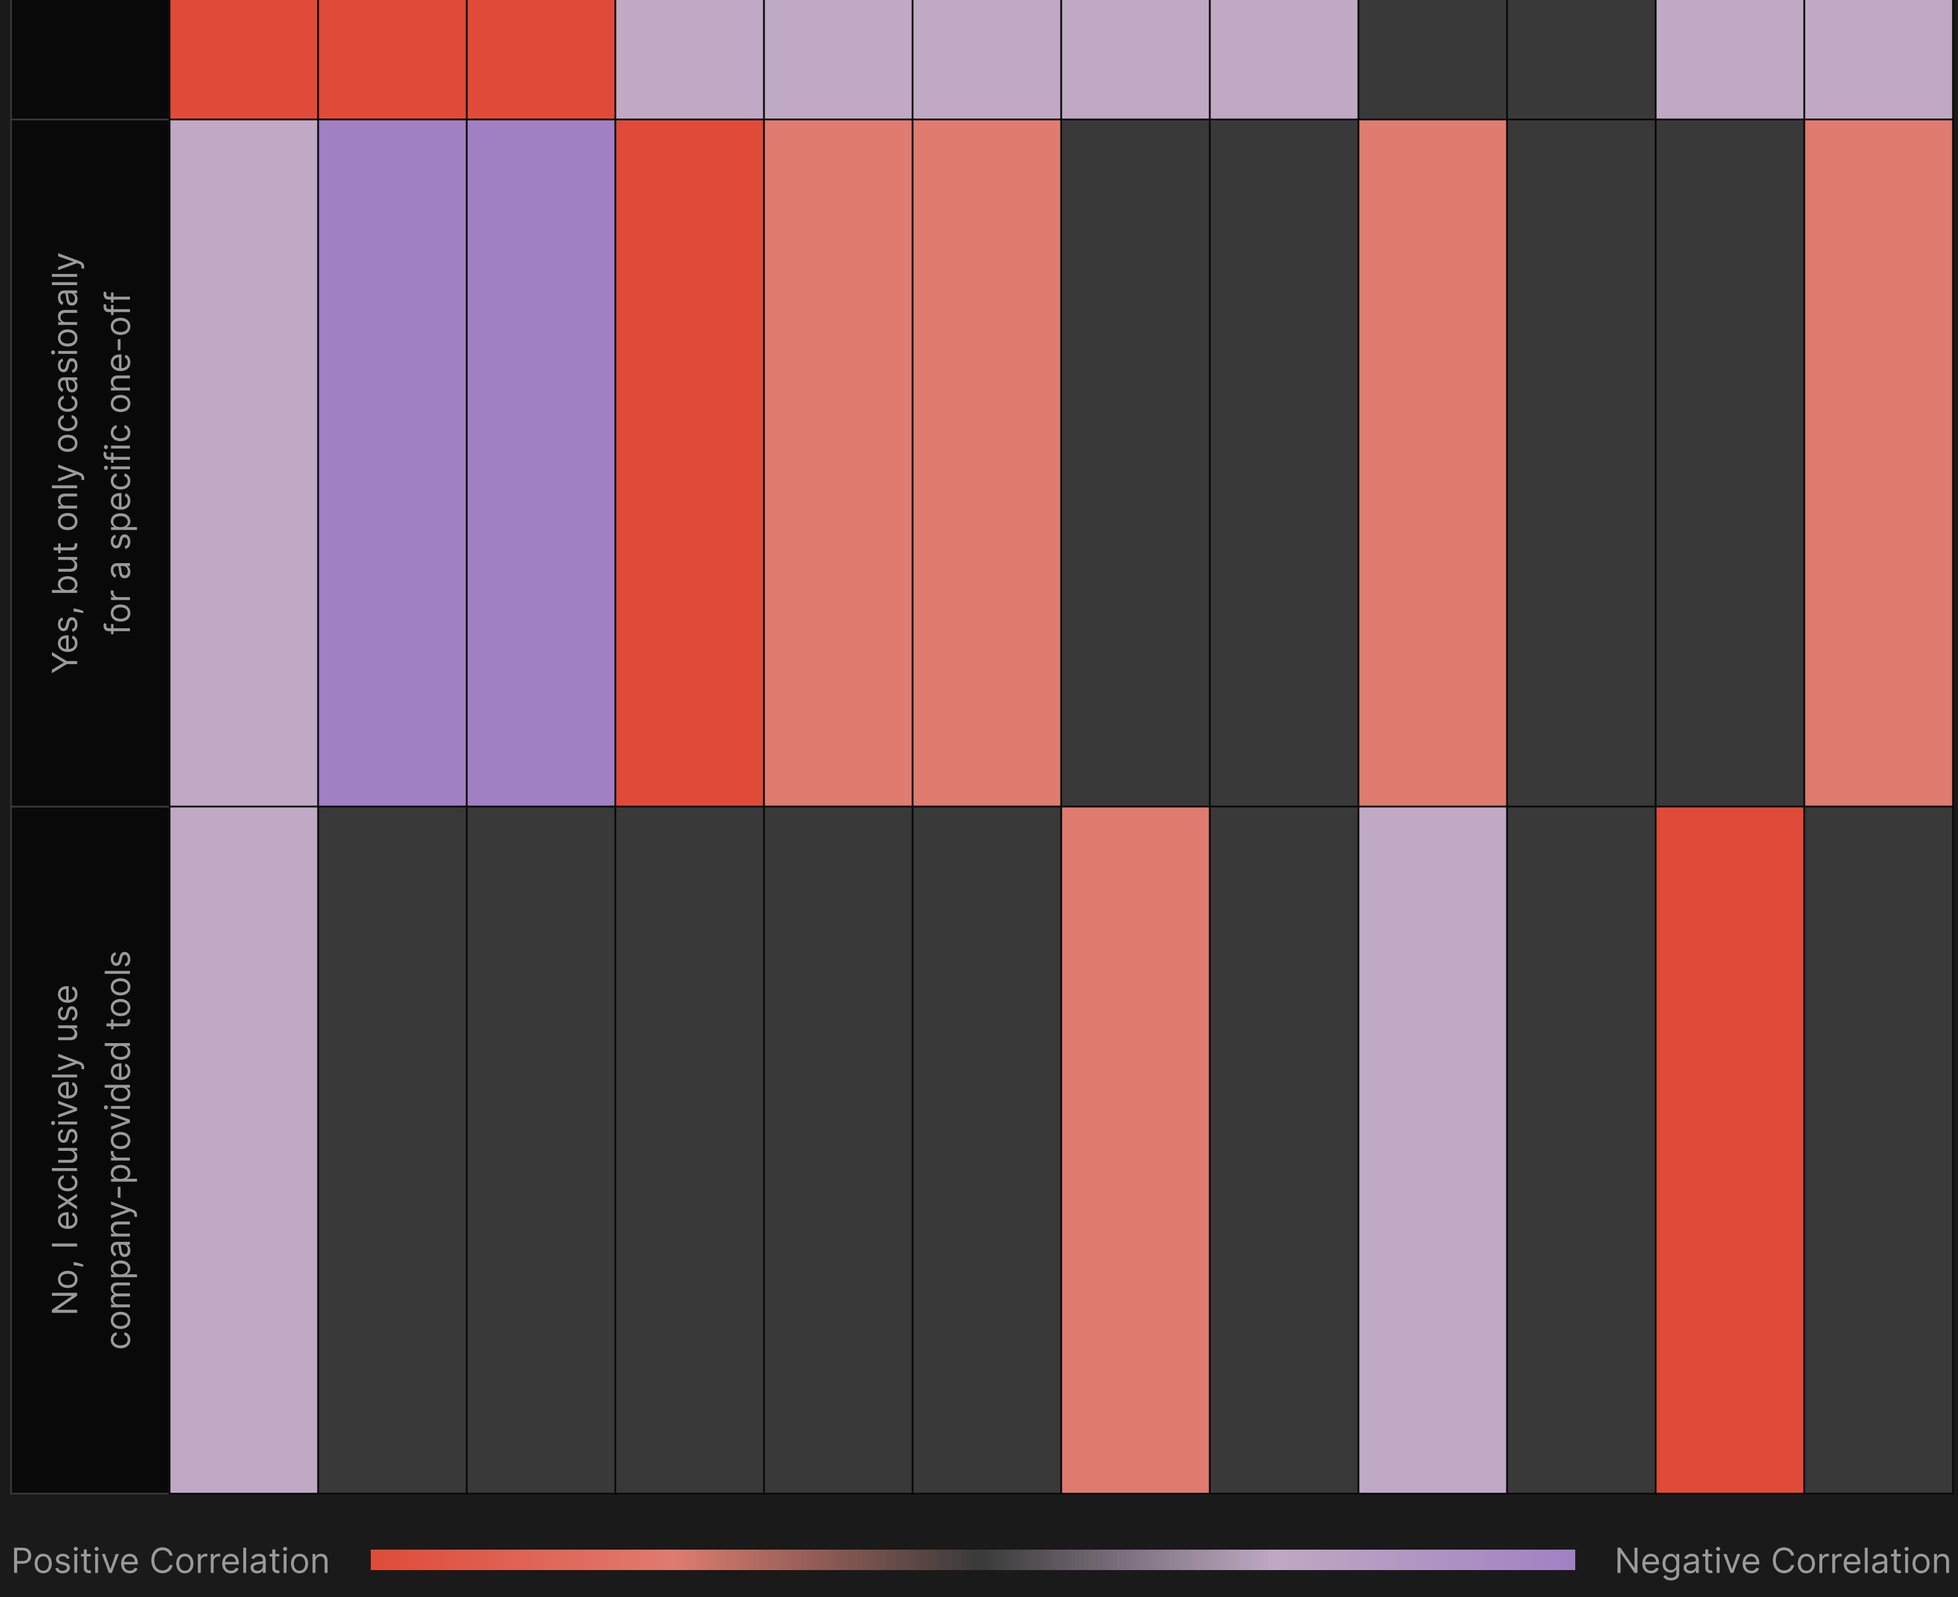

Keeping pace in a system that’s falling behind.

Traditional employee awareness campaigns are not keeping up with the problems of shadow AI.

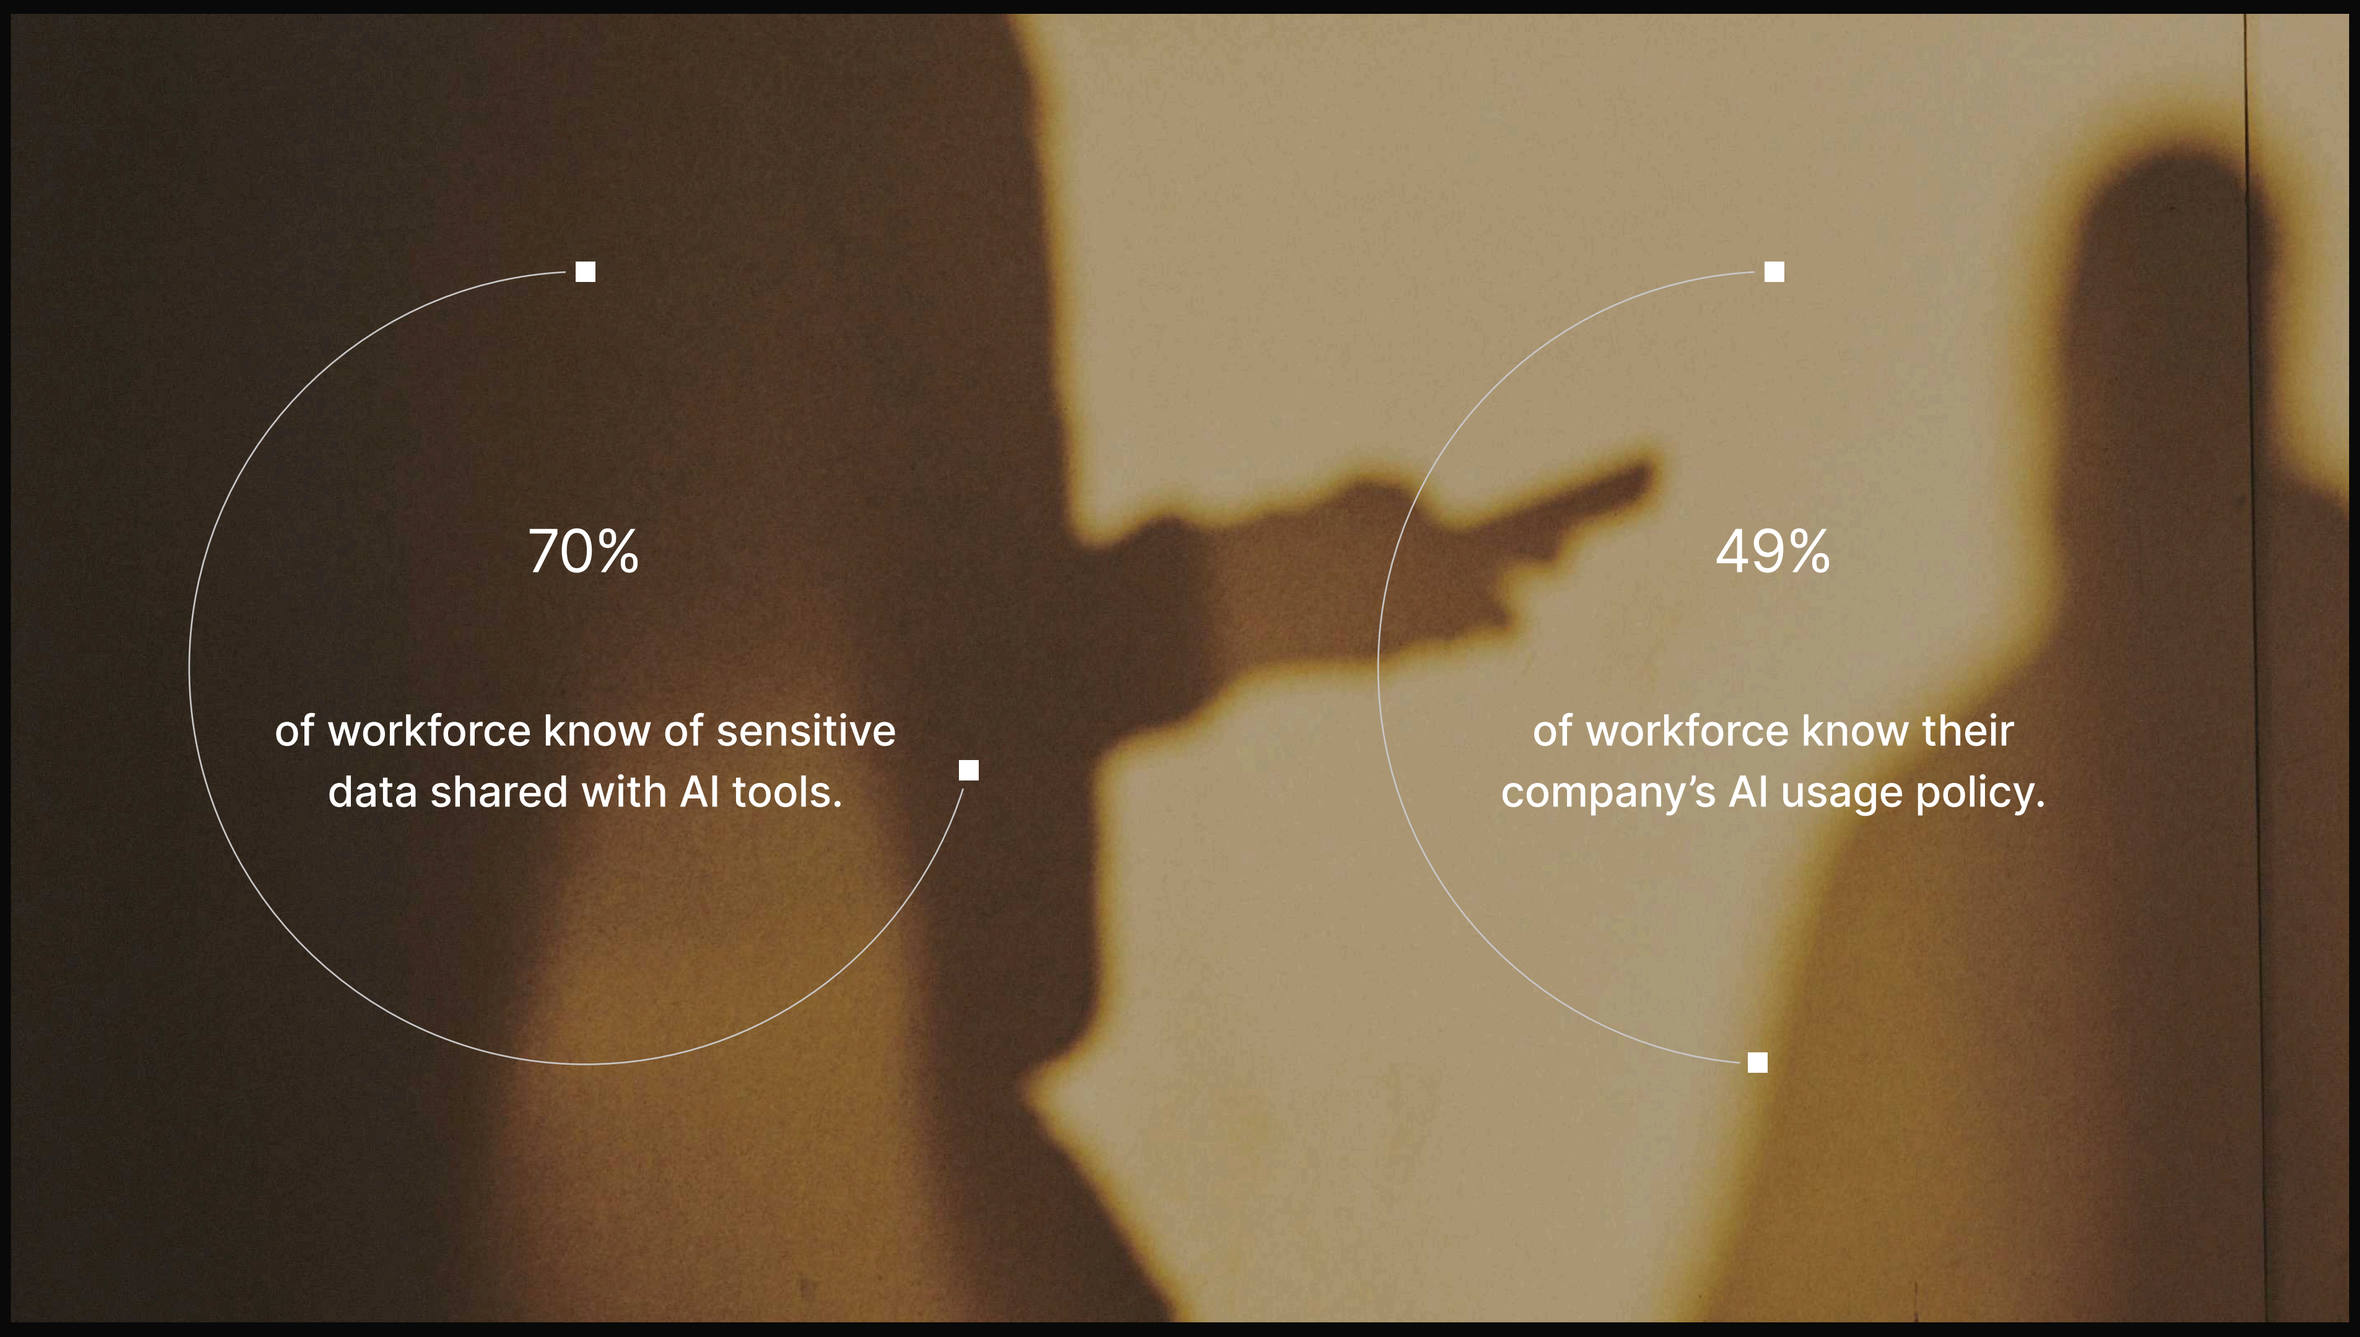

One slice of the workforce have become AI power users, well-versed in the risks of AI and, ironically, the most common source of shadow AI. Many more workers have not yet made it that far in the AI adoption journey. These are not engaged, security conscious participants in AI governance—they are simply unaware of AI security at all.

At best, half of employees can recall their company’s AI usage policy. About one third have received and can recall training on how models work and how to use them safely. These numbers represent both the successes of security organizations in quickly ramping up AI governance, and the inherent challenges in driving large-scale security culture changes.

The results have been middling both in terms of employee awareness and its impact on behavior. Despite the gains in employee awareness of security practices, both employees and security leaders are well aware of sensitive data inappropriately shared with AI applications. Indeed, more employees are aware of sensitive data shared with AI apps in their workplace than know of their AI usage policies.

Workforce Awareness of AI Policy

Just under half of employees rate themselves familiar with the contents of their AI usage policy. A third of employees don’t believe their employers have one.

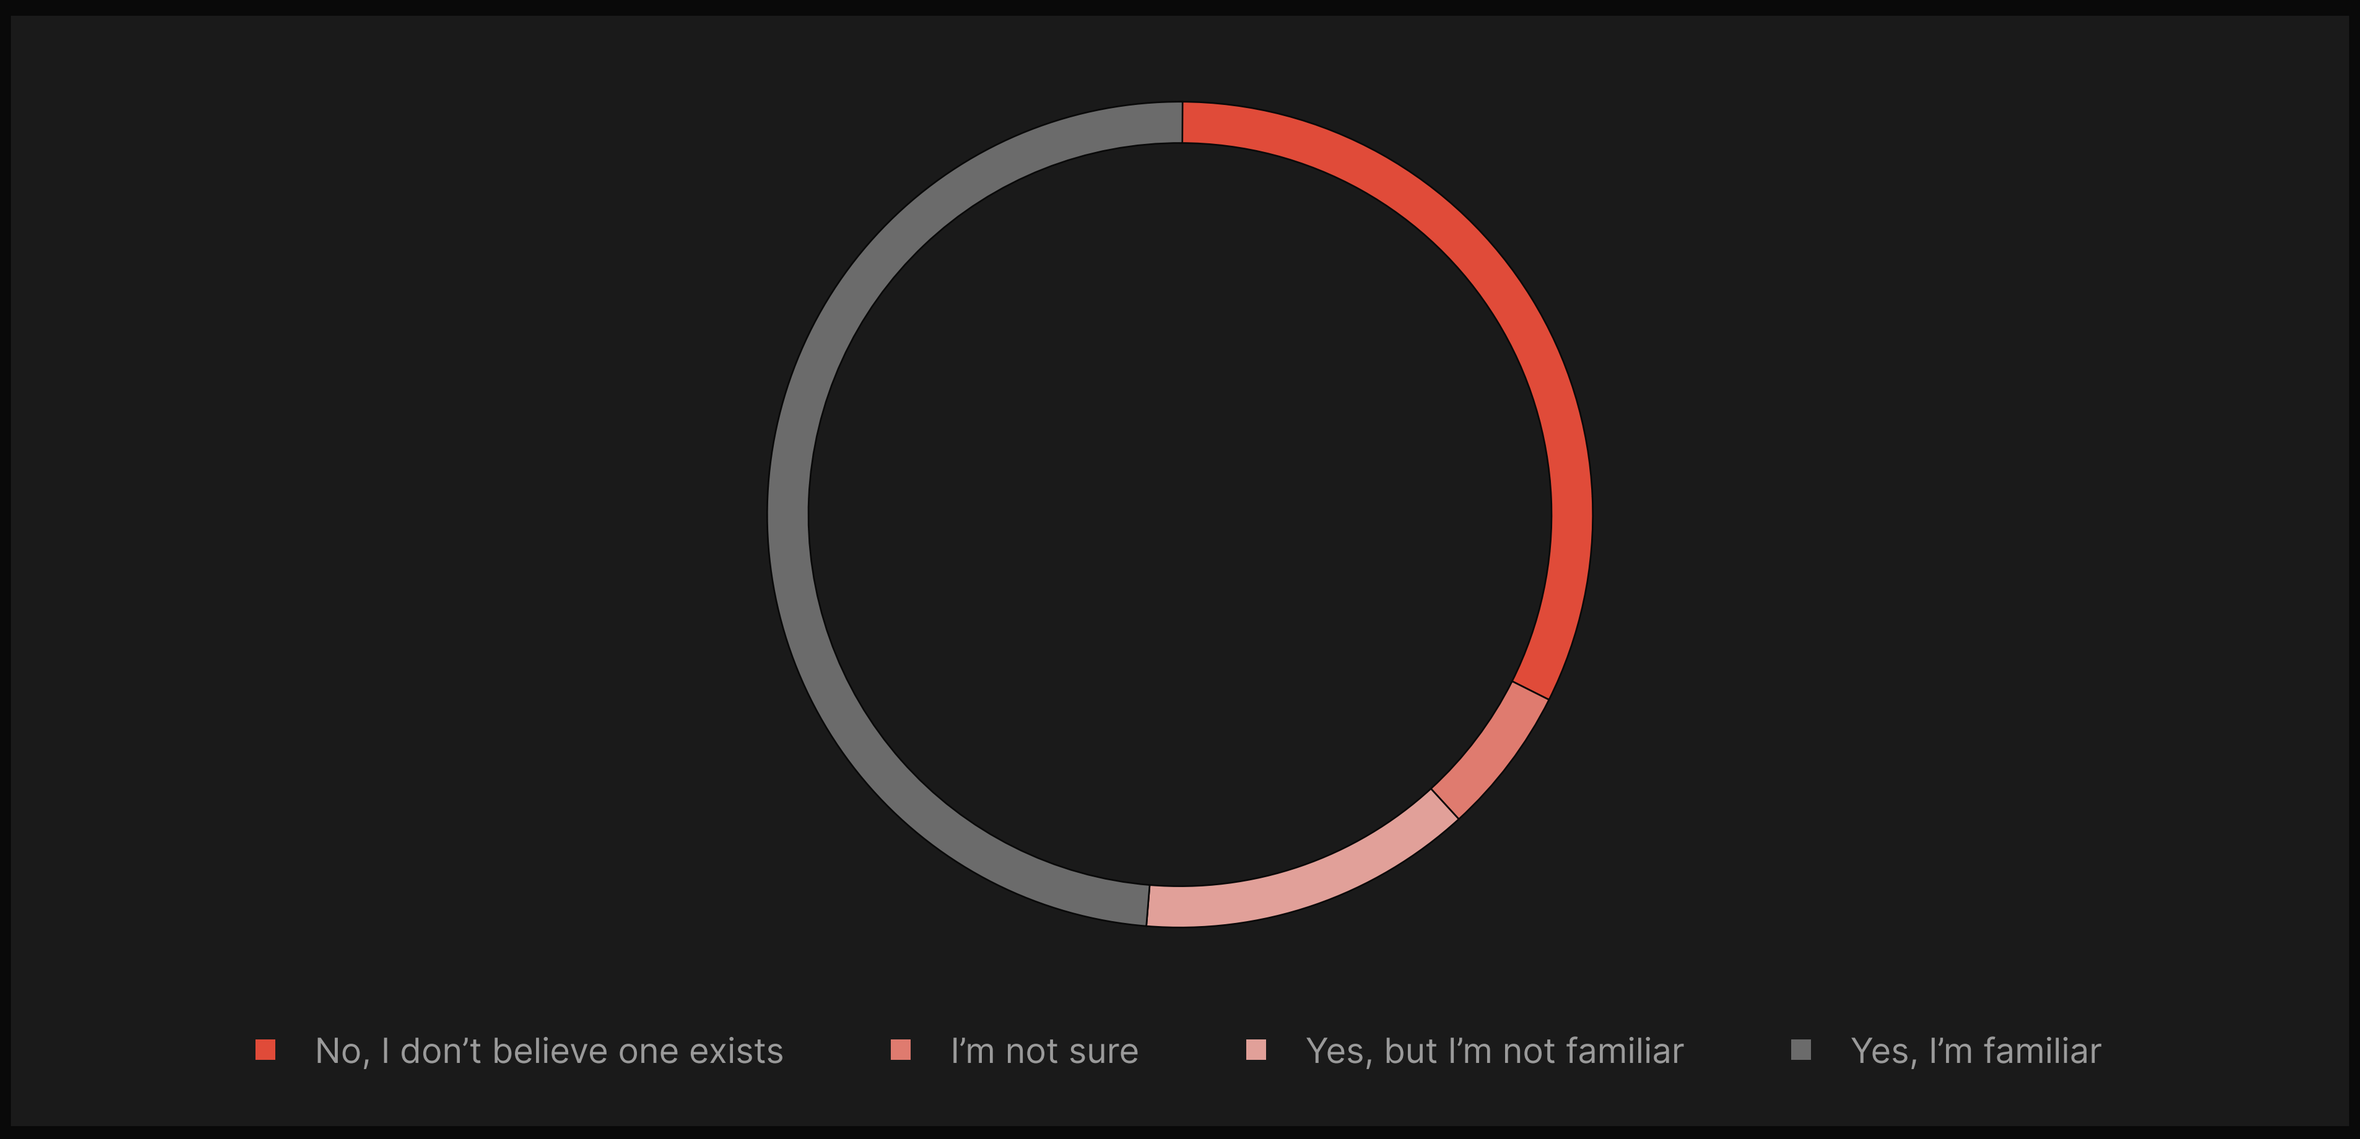

Workforce AI Security Training Retention

A third of employees recall the contents of AI safety training. About half have not received any training on using AI securely.

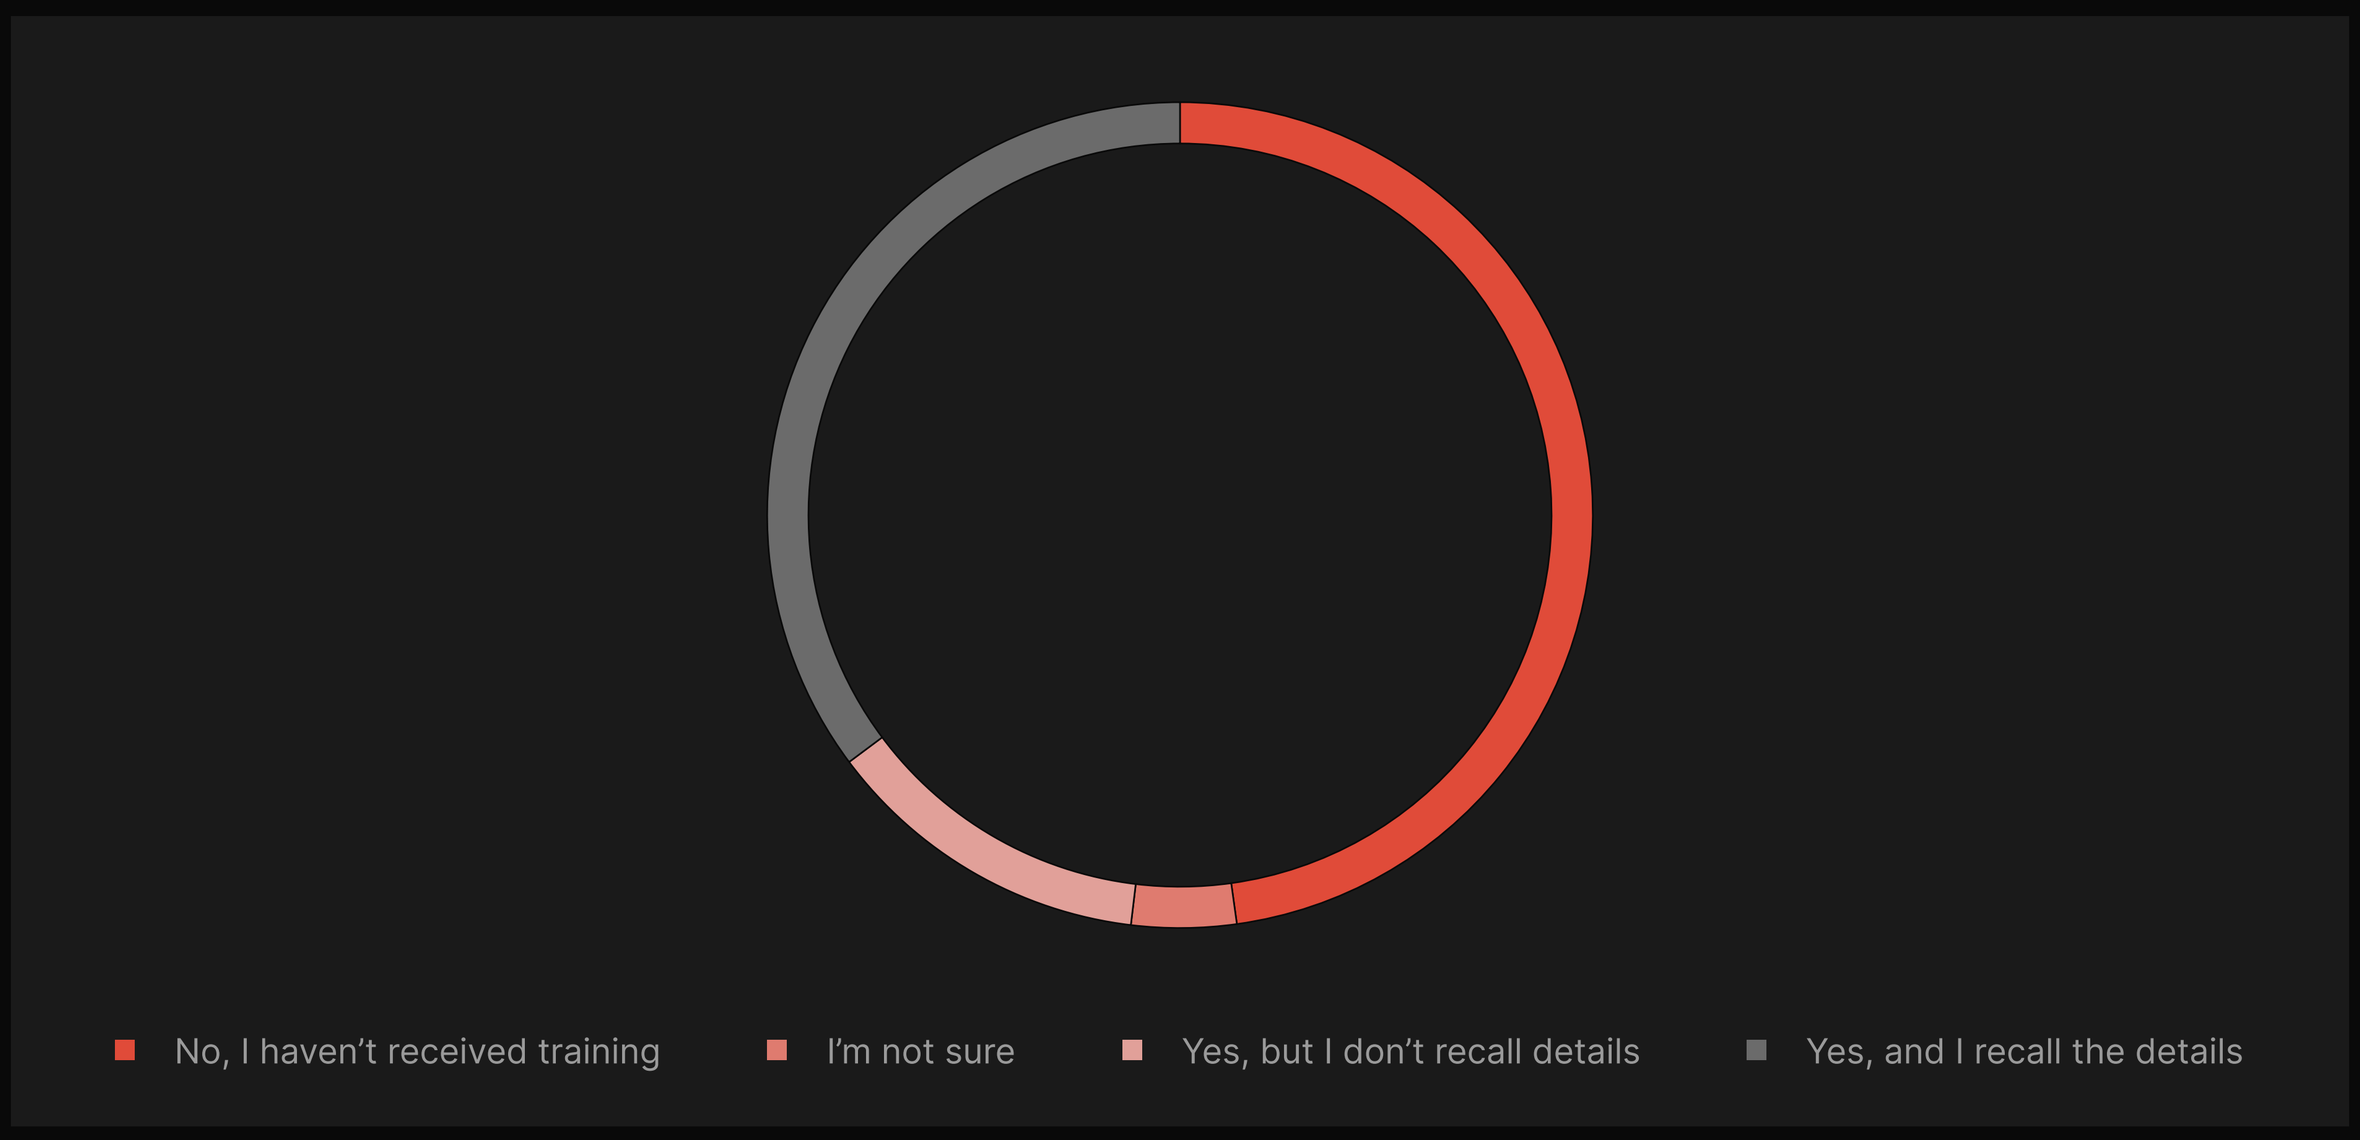

Workforce Understanding of AI Tool Inputs

About one third of employees are confident they know how their tools use inputs; another third feel sure about some tools but not all. One in eight employees don’t know and don’t care.

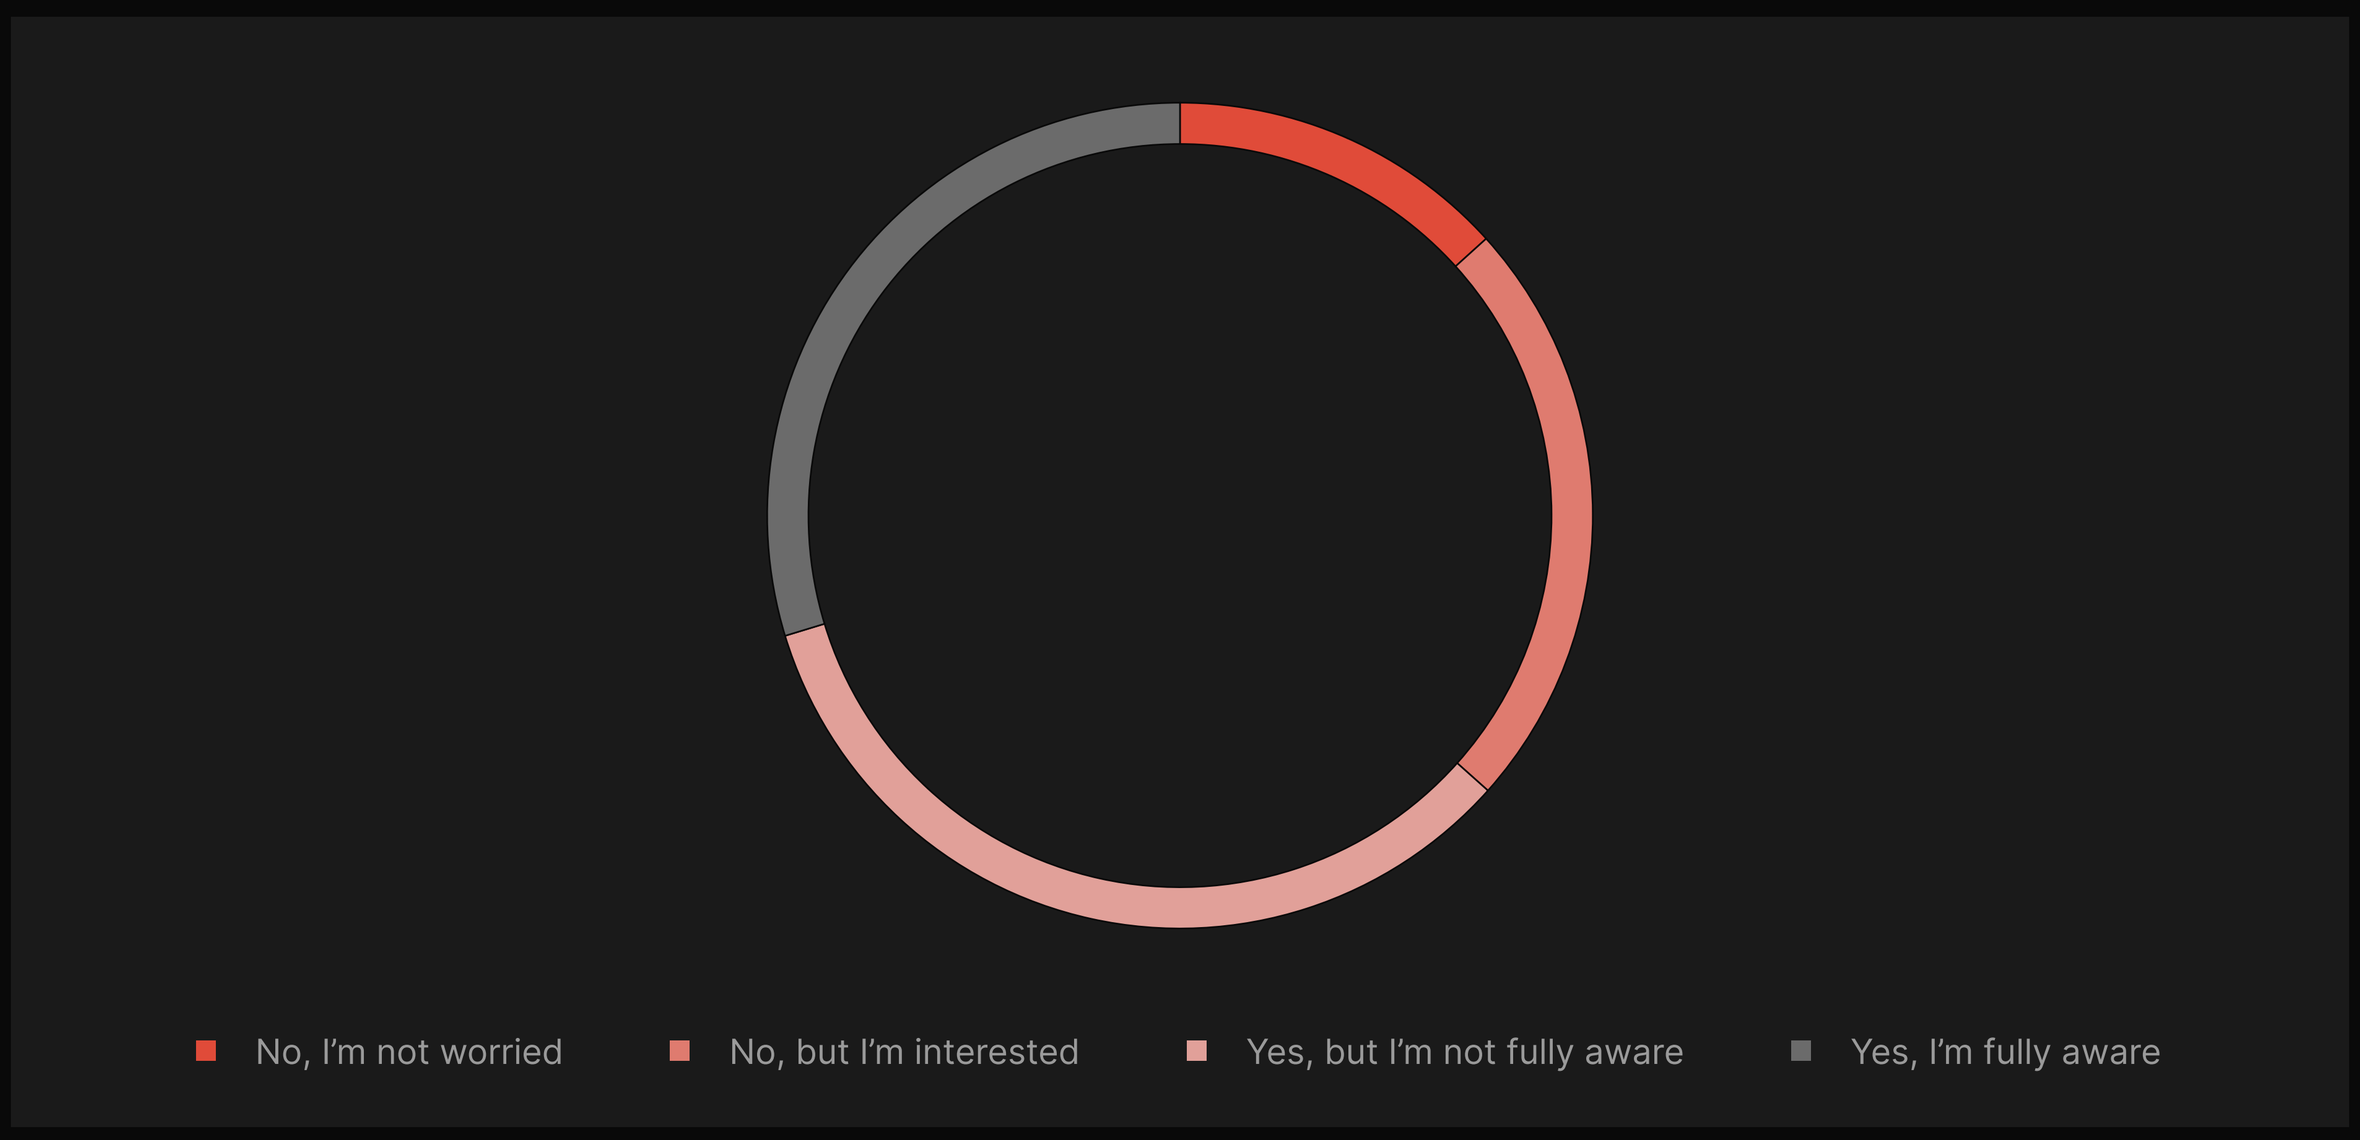

The State of AI Security Awareness

A clear policy alone doesn’t guarantee secure behavior. While many employees are aware of AI guidelines, most lack consistent training or a firm grasp of how AI tools handle inputs, leaving the majority in need of further support.

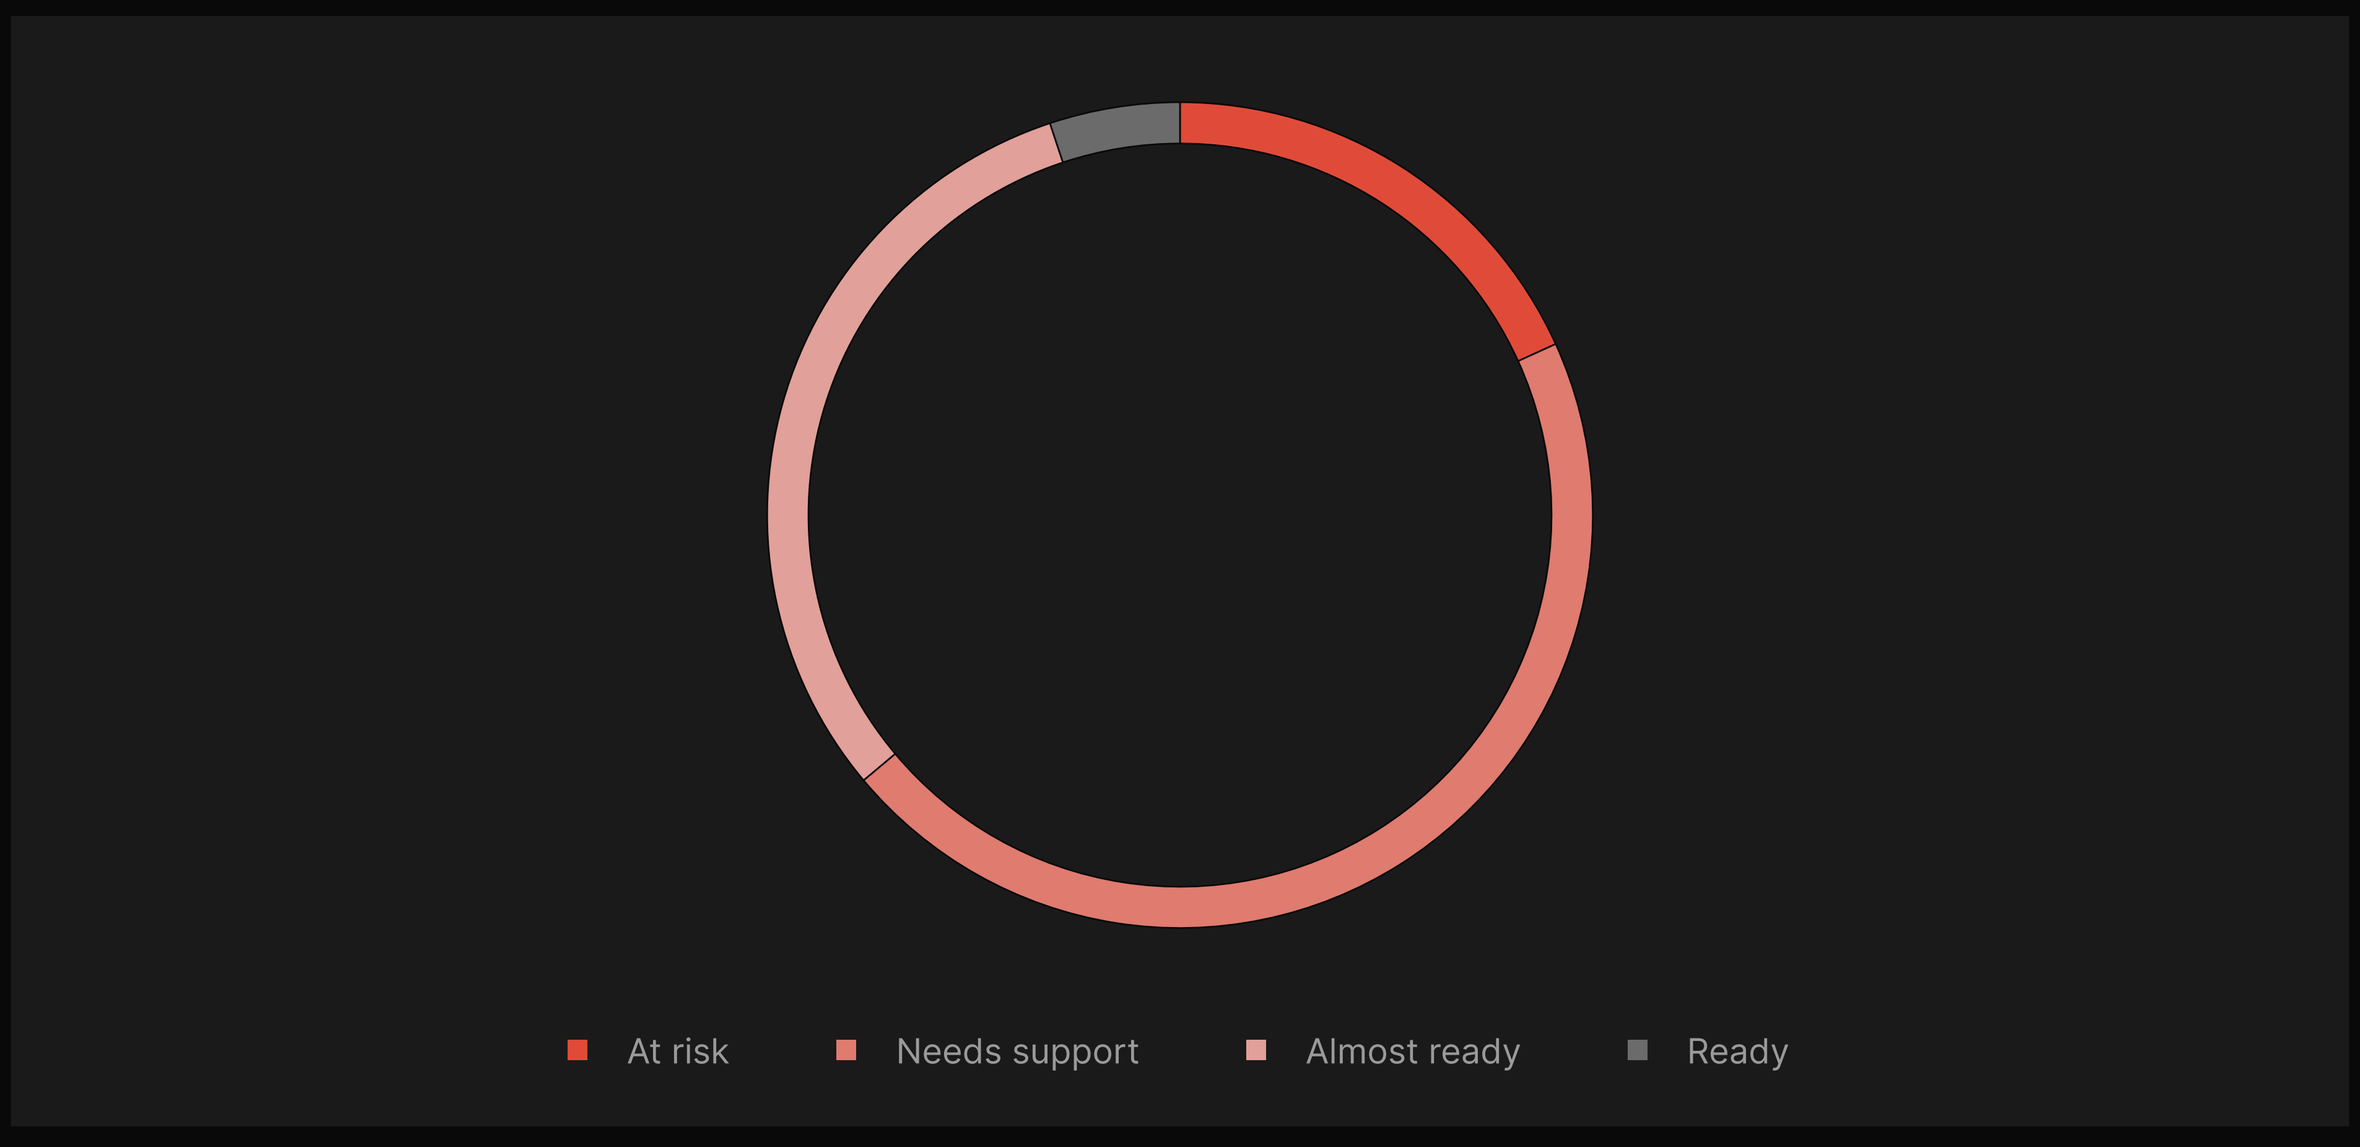

Knowledge of AI Data Leaks

Both security leaders and the general workforce know of sensitive data shared with AI tools. Security leaders have greater awareness of AI data leaks than the workforce, as they should—-they just don’t have a way to stop them.

Knowledge of AI Data Leak Types

When asked whether they were aware of different types of confidential data being shared with AI tools, workers and security personnel had similar experiences—though the higher rate amongst the workforce suggest some leaks might be undetected.

AI is replacing trust between people.

A quarter of workers trust AI more than their managers and co-workers, eroding the foundation of workplace governance.

Shadow AI is, at minimum, a crack in the relationship between leadership and workforce. But just how deep does it go?

To measure how much AI has shifted trust in the workplace, we asked employees who they trust most for information: their manager, their co-workers, search engines, or their favorite AI tools. Despite very public cases of AI hallucinations, trust in AI often ranks at the top of the stack.

Here, some differences emerged between workers in the UK and US, where 20% and 28% of workers selected AI tools as their most trusted source of information, respectively. That difference might reflect broader differences in the cultural operating environment, including the UK's willingness to regulate AI compared to the US's reluctance.

The effect of the operating environment on risk appetite can also be seen when slicing the data by industry. Employees in Technology, Finance, and Manufacturing are as likely to trust AI as they are to trust their manager, whereas workers in Healthcare—where the stakes of liability are far higher—greatly favor their manager.

Not surprisingly, users who trust AI more than people also use shadow AI more often than their peers. While the rates are about the same for employees who trust any other source, those who trust AI mostly use their own AI tools.

Furthermore, we can see the impact of a degraded trust environment in the open secret of AI data leaks. Security leaders and employees are both well aware that sensitive data is being sent to AI tools, normalizing an environment of disregard for governance.

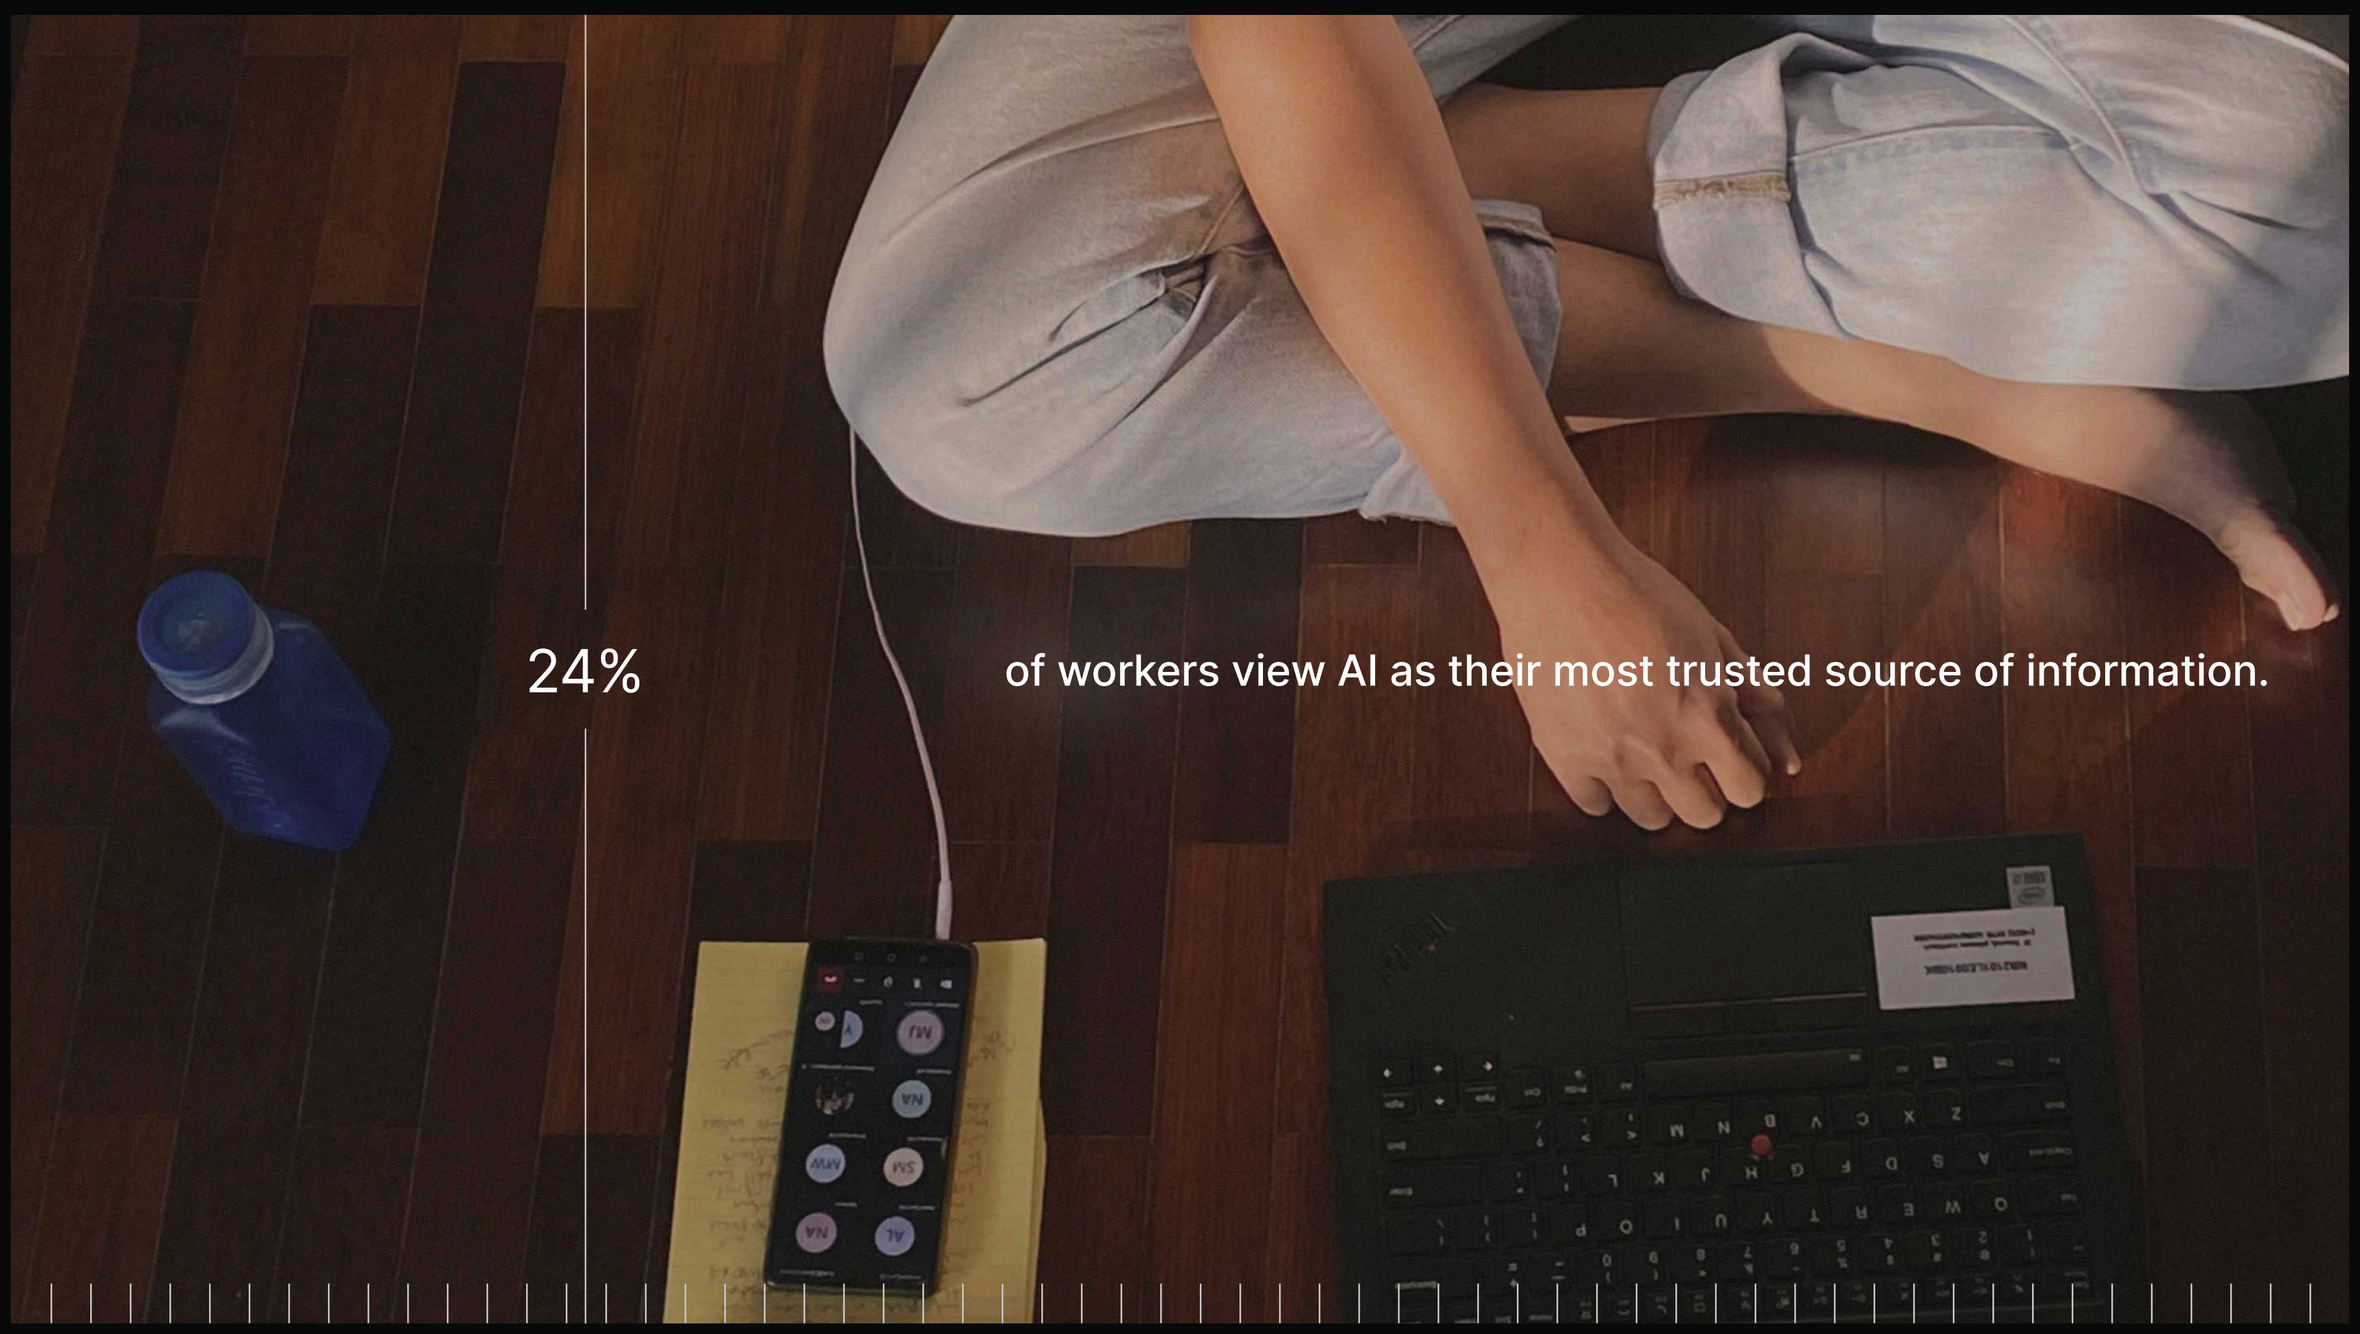

AI vs. Human Trust at Work

One in four workers view AI as their most trusted source of information. That percentage rises as other environmental factors like geographic region influence the culture of AI adoption.

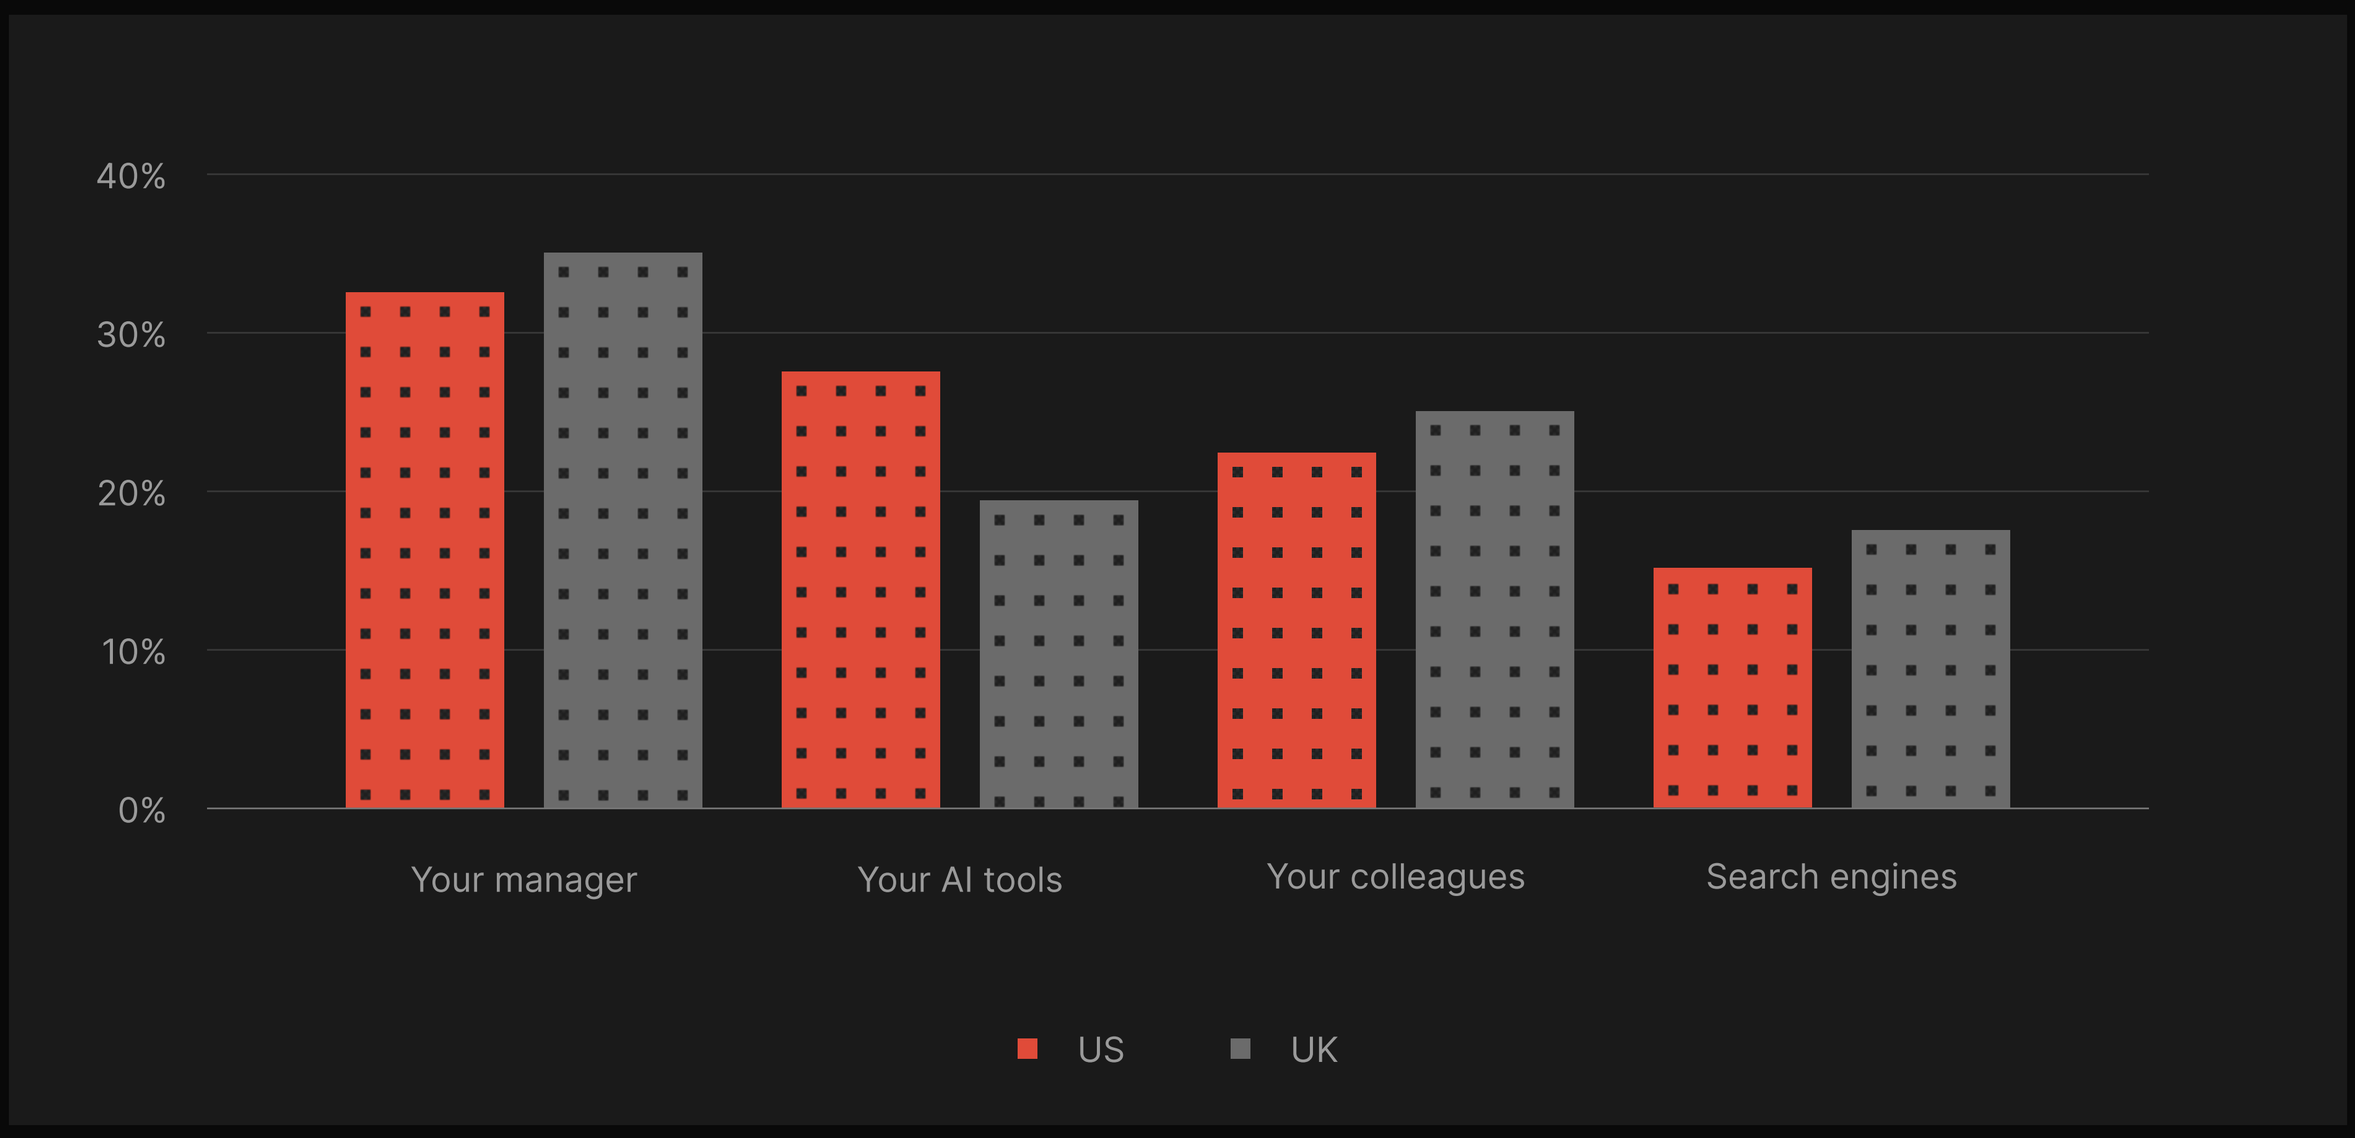

AI Trust by Industry

Industry is another part of the operating environment that influences AI adoption and, in turn, the rate at which it is the most trusted source of information—in some cases surpassing both managers and co-workers.

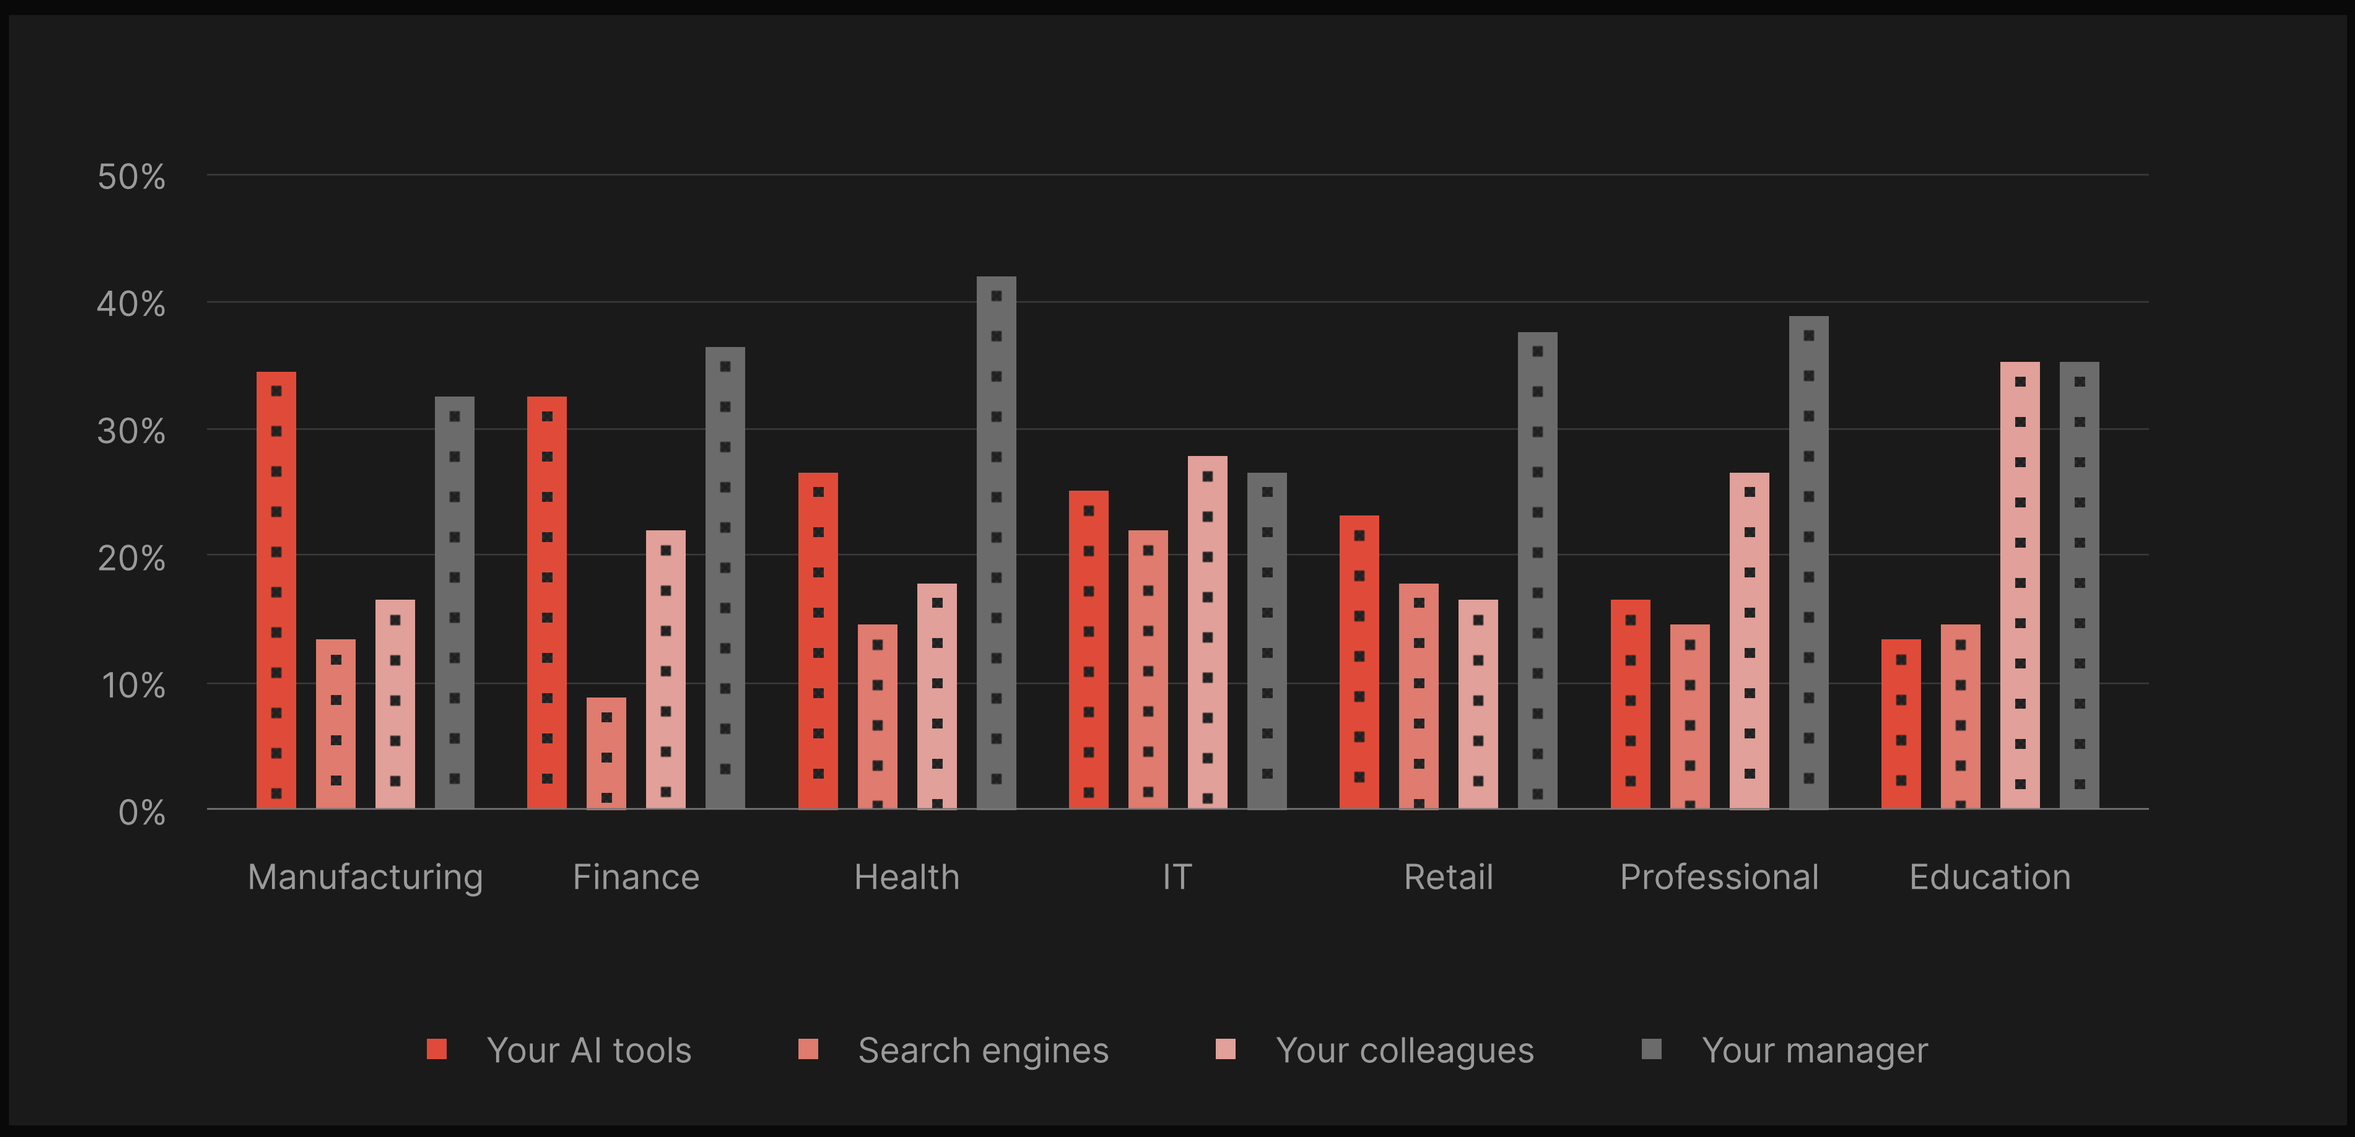

How Workforce Use AI Tools

Employees who view AI tools as their most trusted source of information are far more likely to use shadow AI tools as part of their regular workflow.

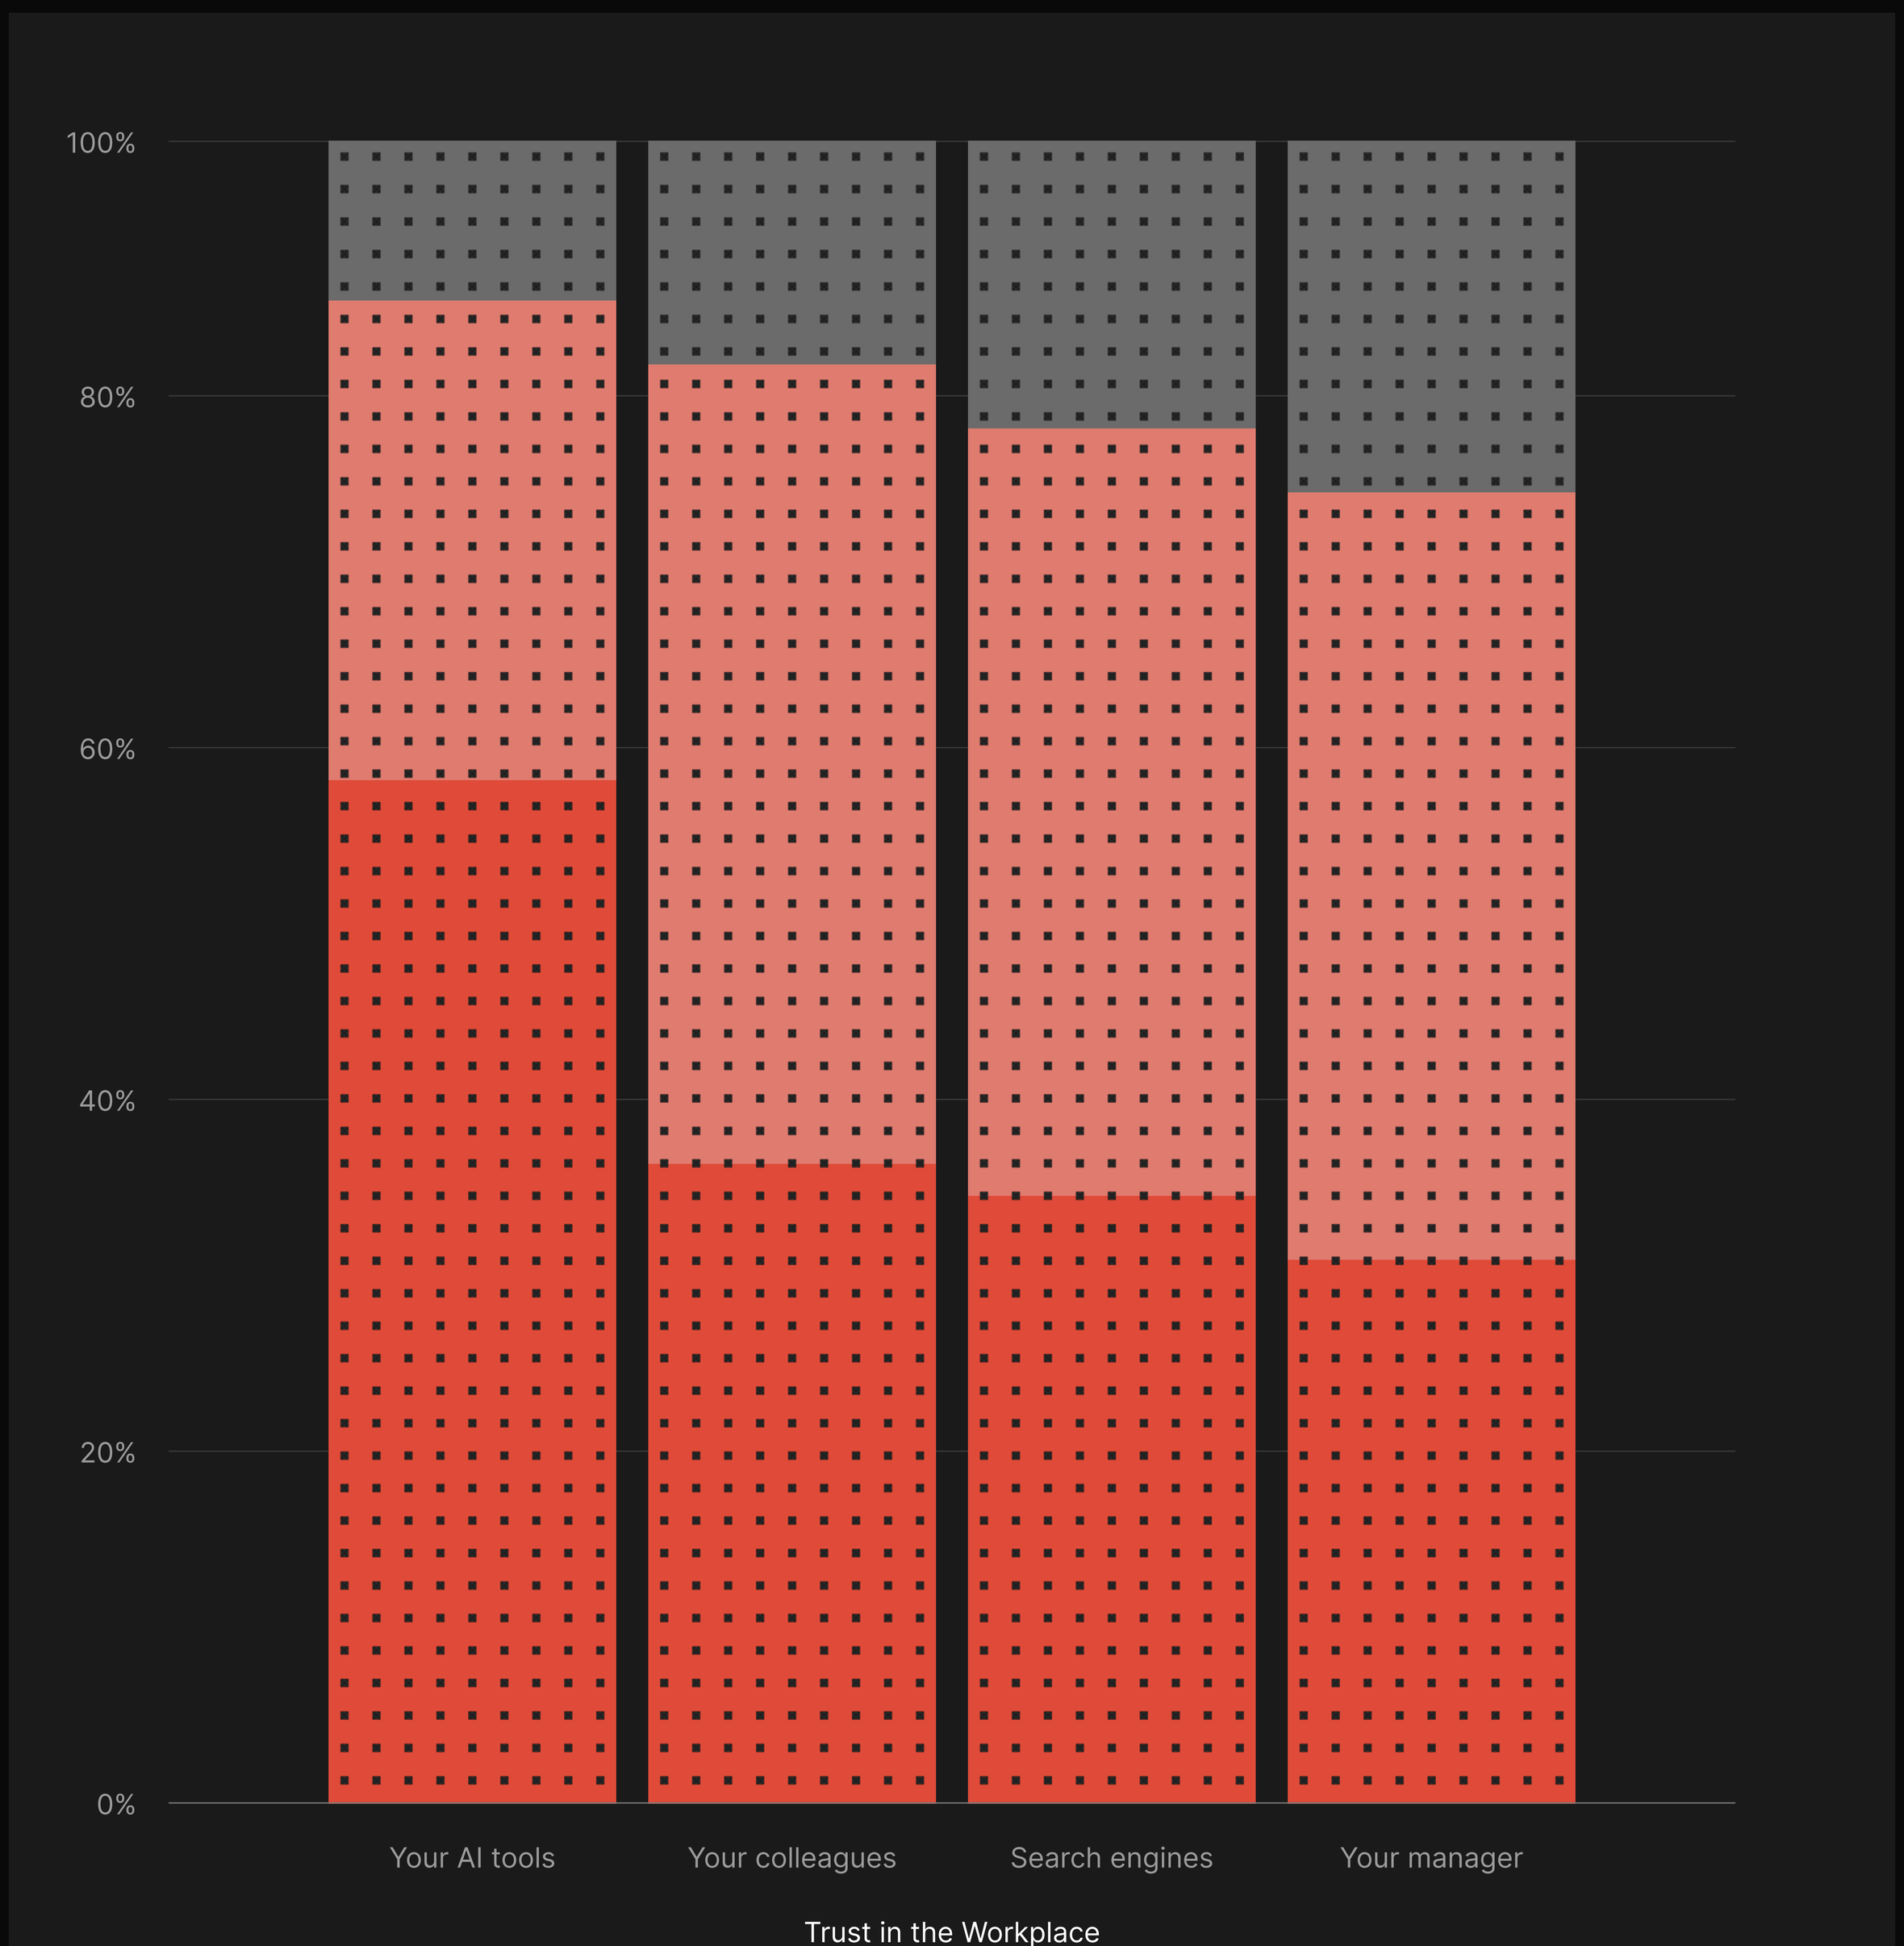

Security's Estimate of Shadow AI

When asked what percentage of their workforce they believed were using shadow AI, security leaders tended to underestimate how common it really is.

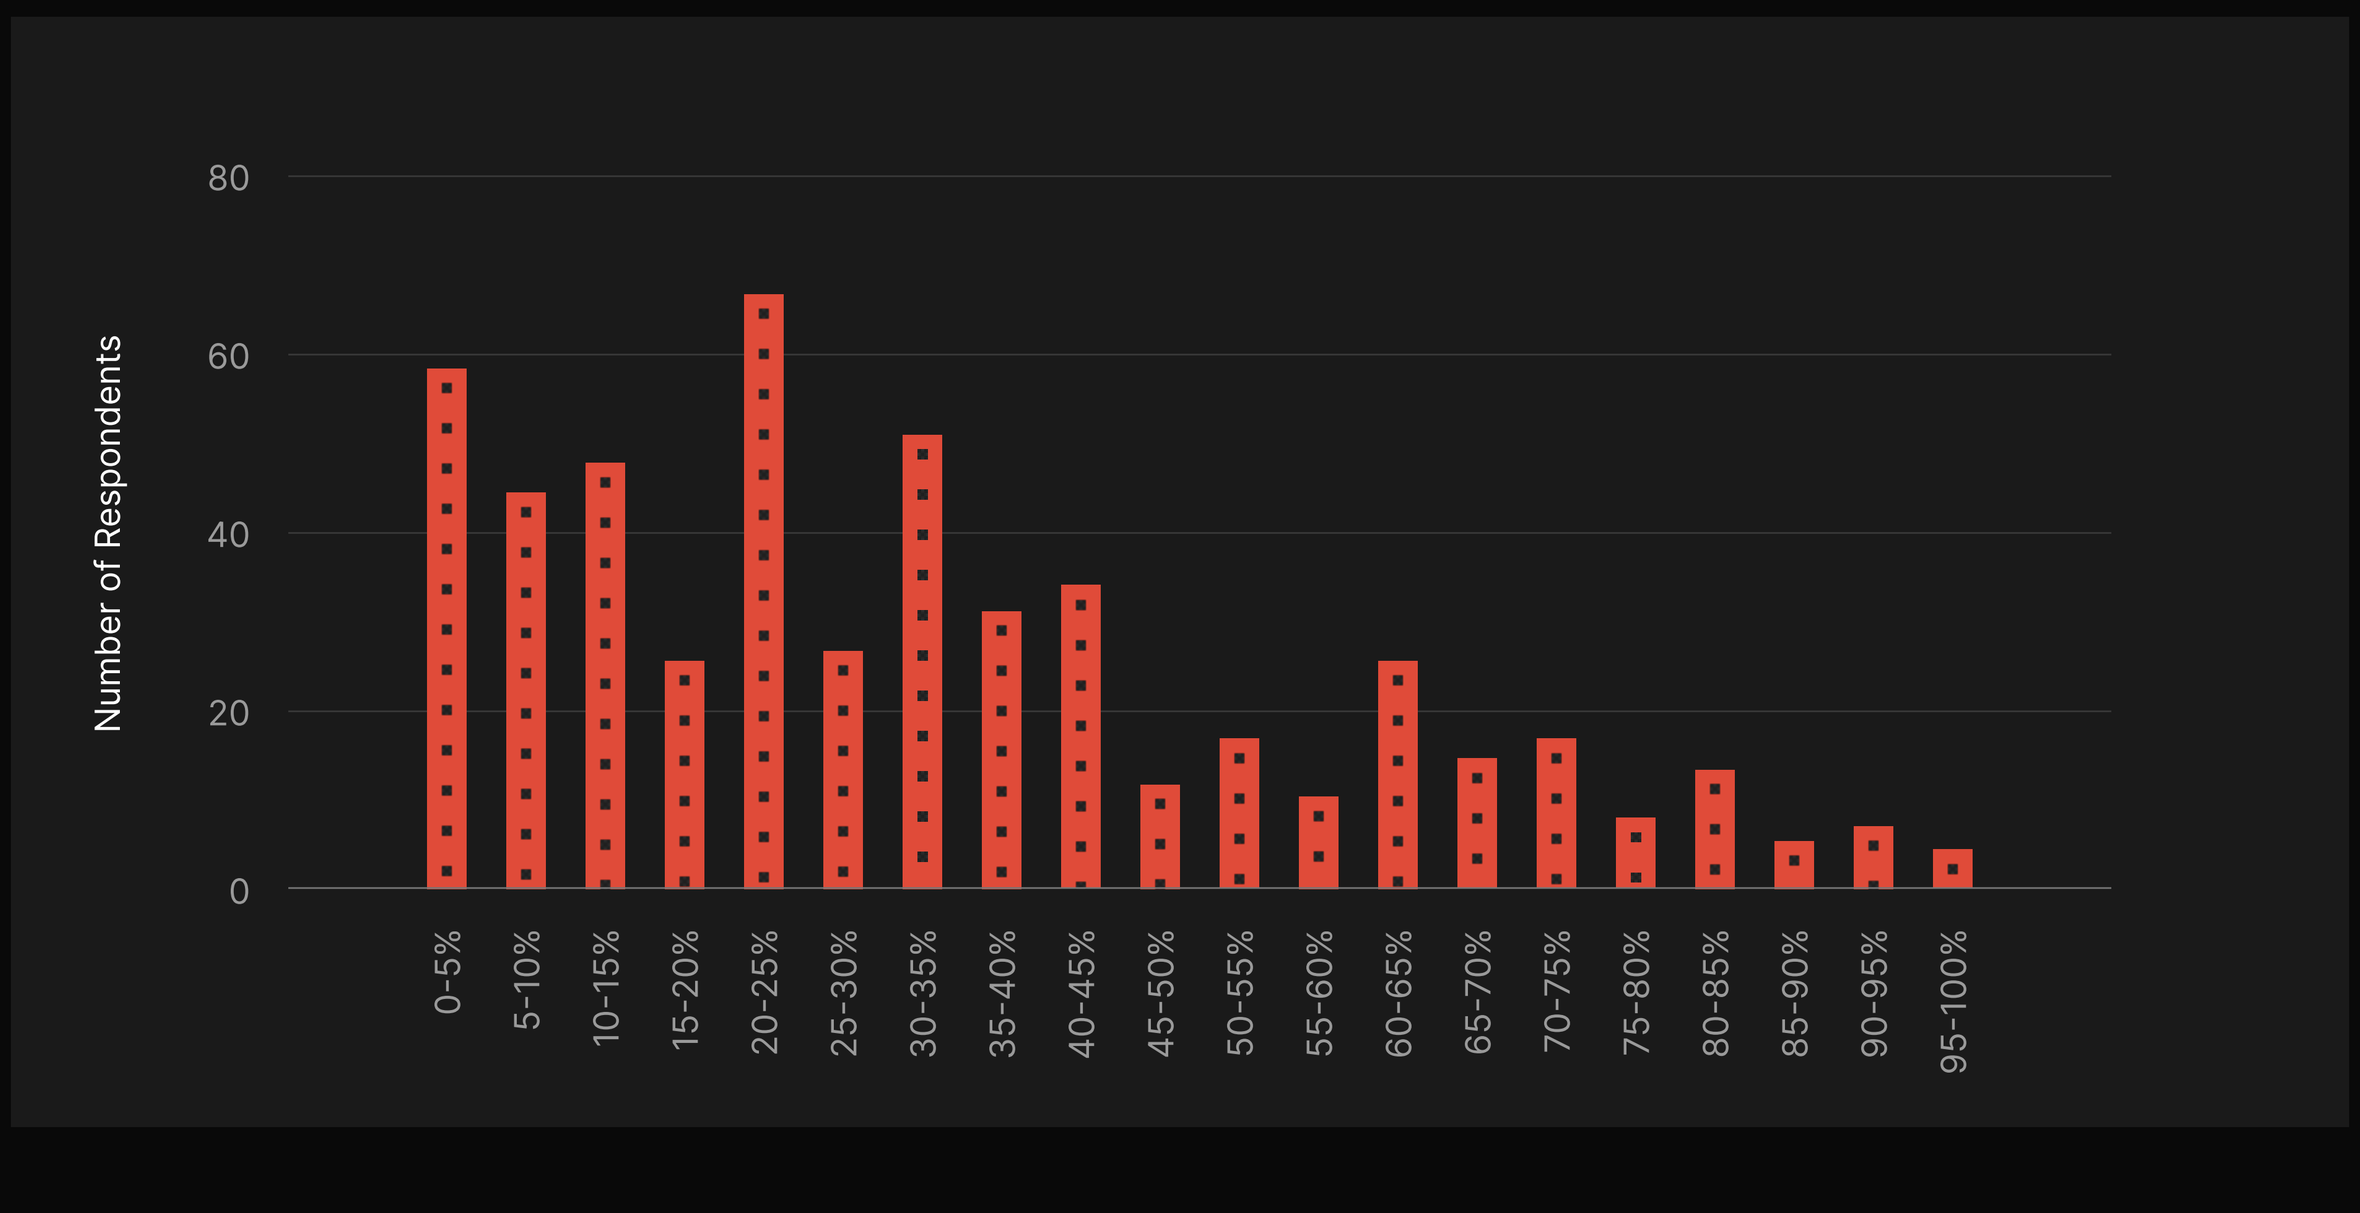

Data Loss is Security's Top Concern

When asked to rank seven potential concerns with AI, security leaders most frequently ranked "employees sharing confidential data with AI tools" as their number one worry.

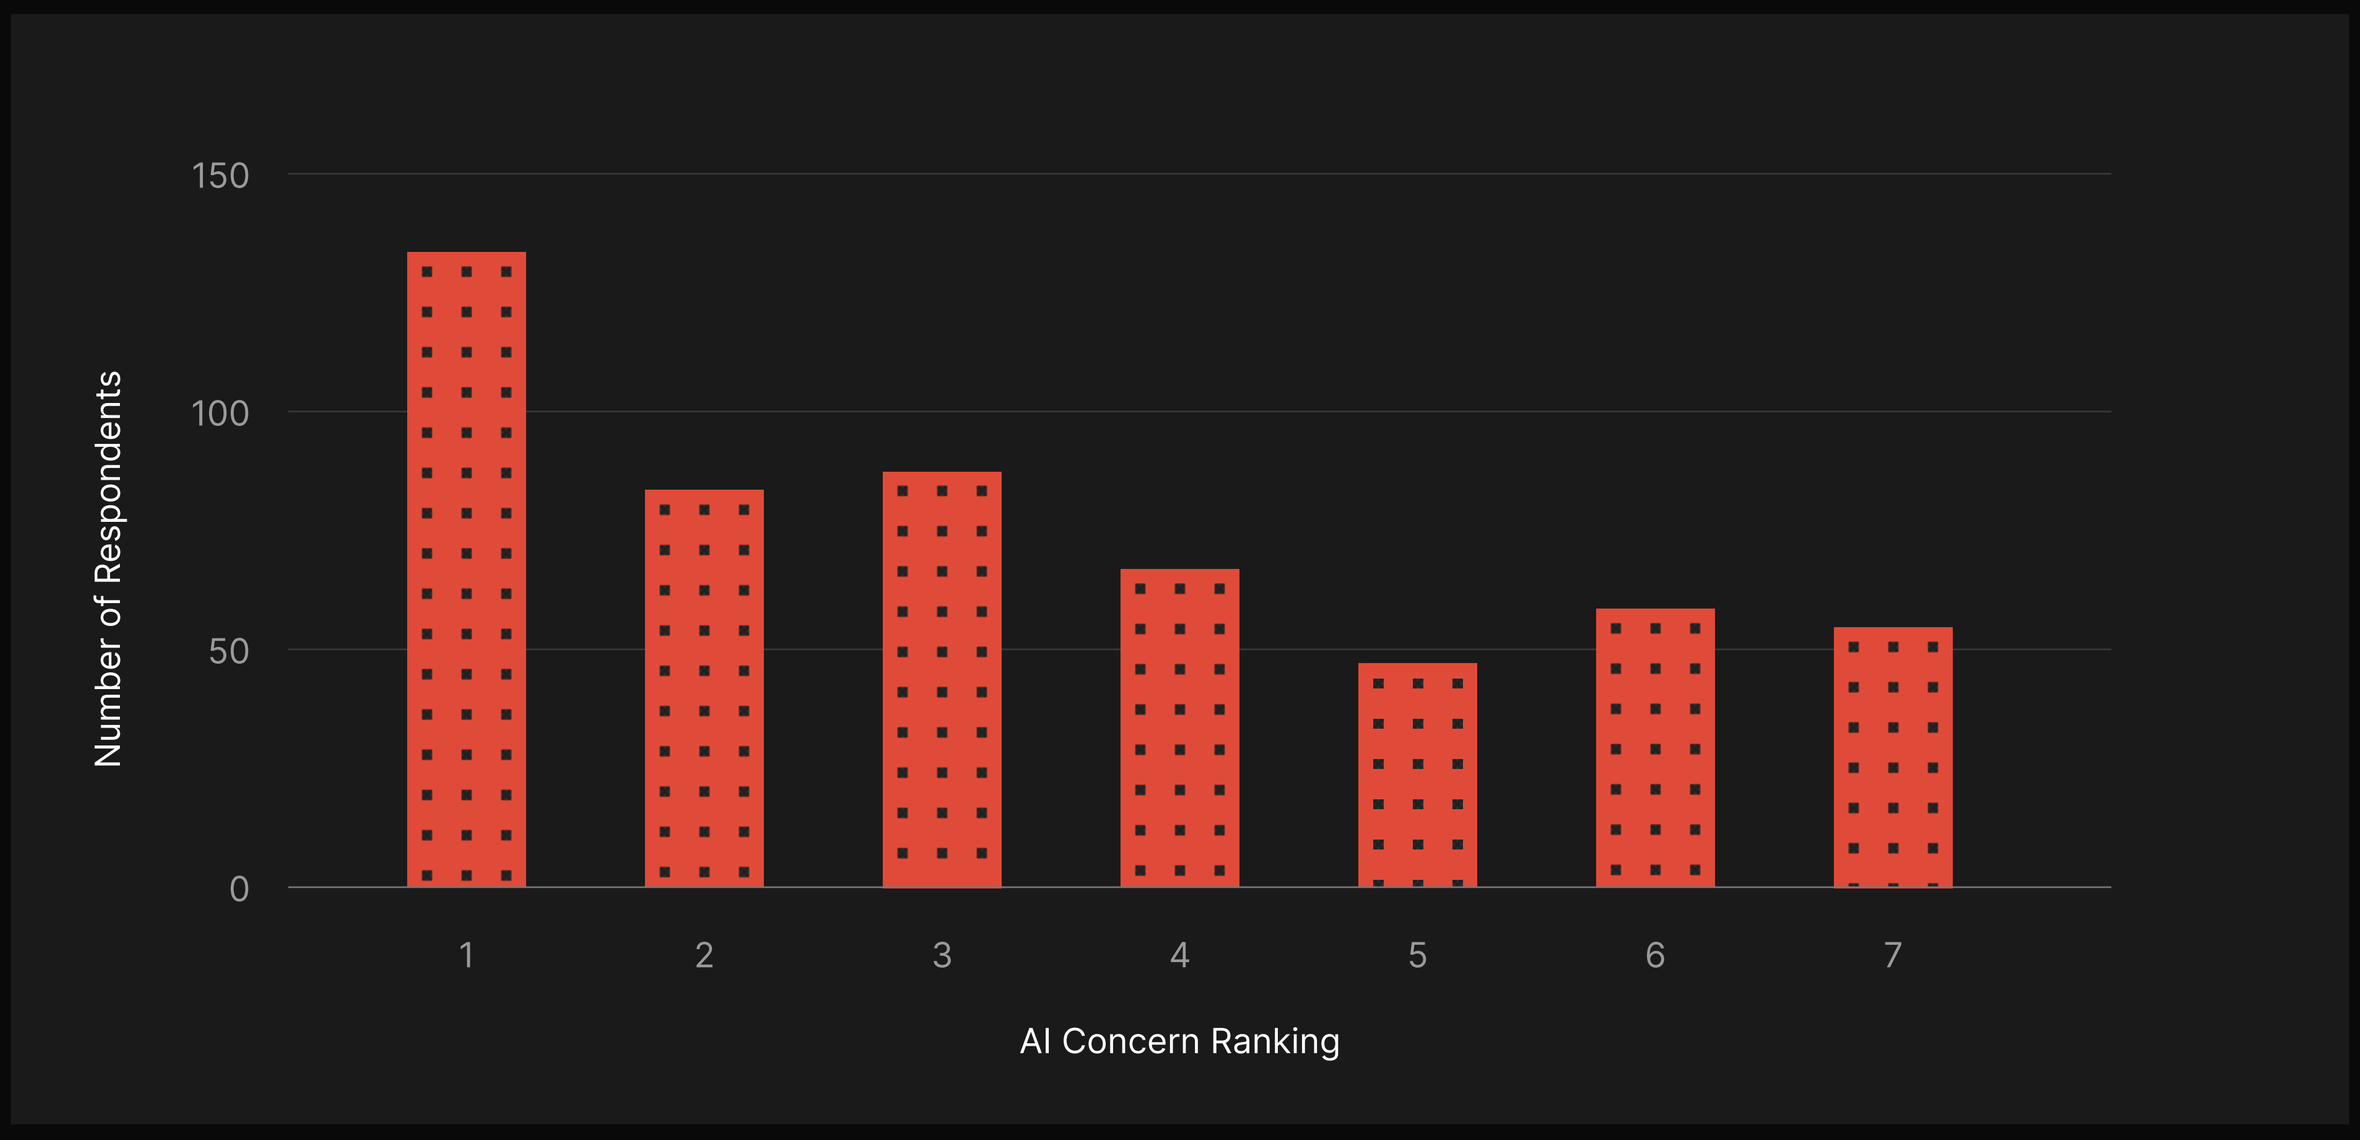

Secure behavior isn’t forced. It’s fostered.

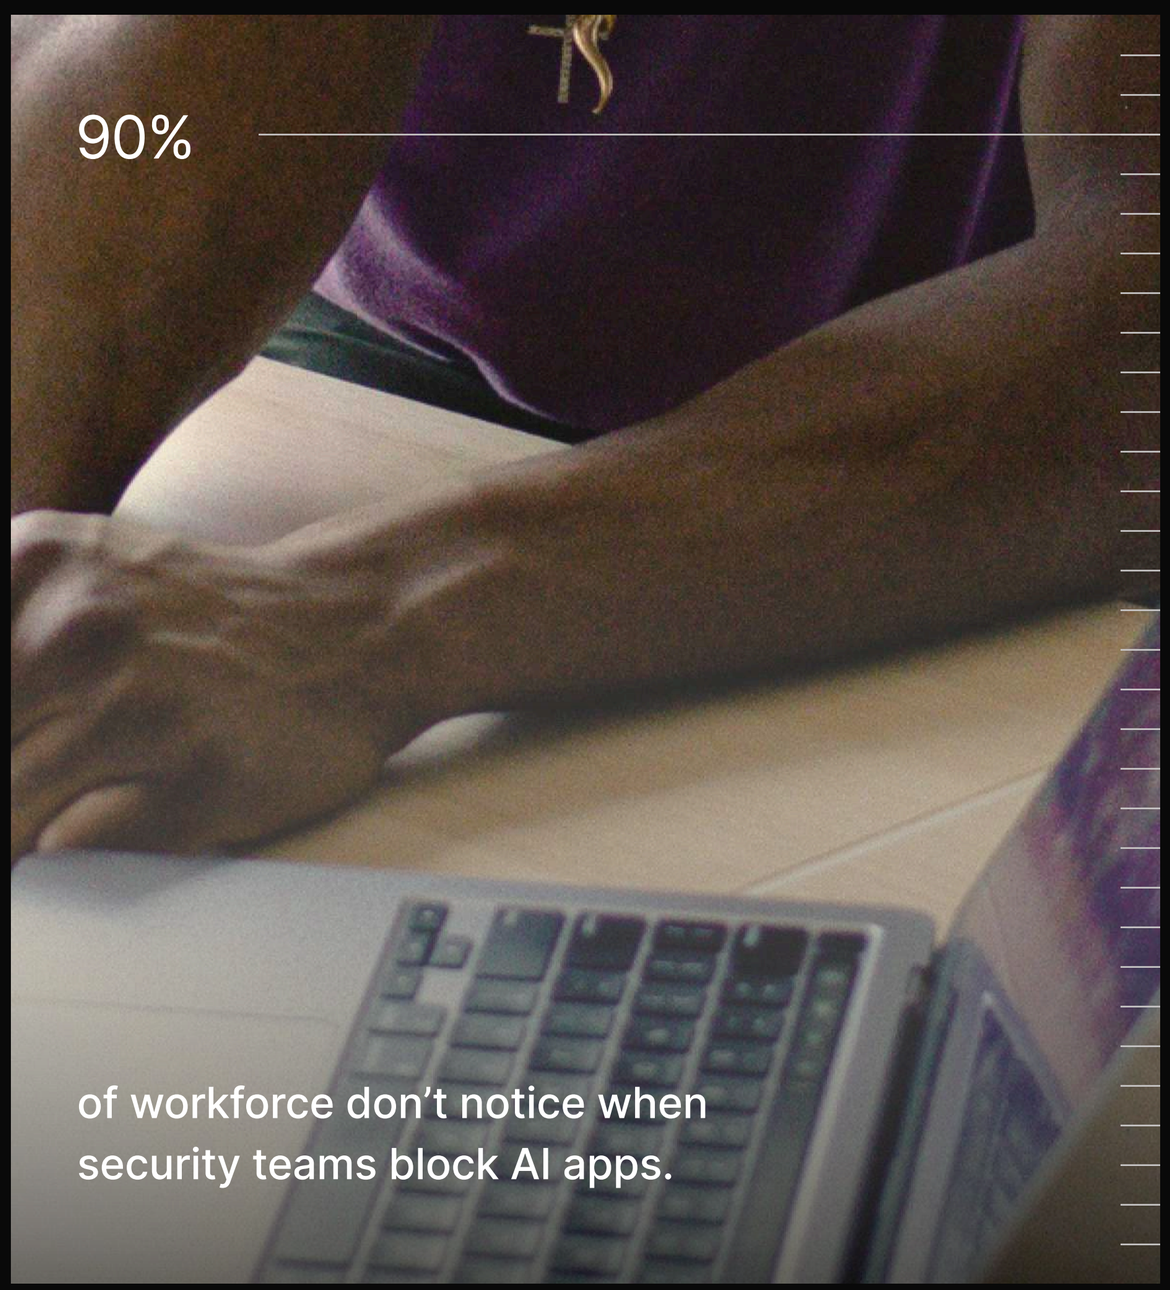

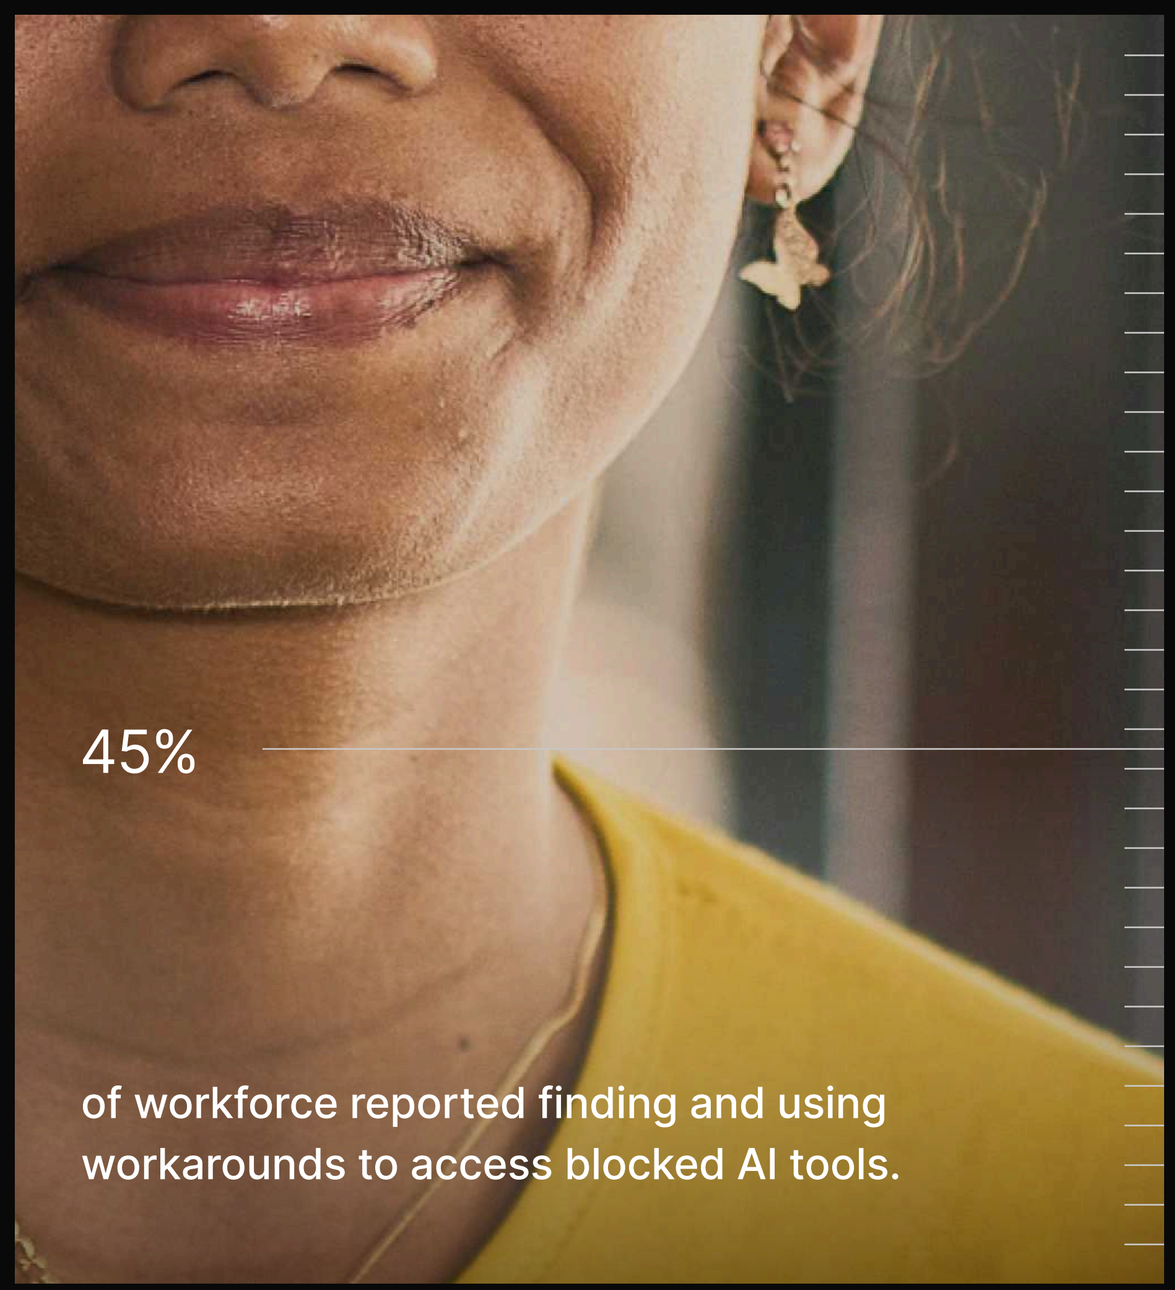

Nearly half of employees find ways around AI restrictions, undermining policy enforcement and visibility.

As employees gain confidence with AI, particularly through security awareness training, they also tend to use more unapproved AI tools. The solution to the shadow AI problem is certainly not to reduce employee training; instead, we have to find a way forward.

One method is to take a firmer stance. Instead of asking employees to follow AI policies, employers can block network access to unapproved tools. This method appears only partially effective. First, we found that security leaders were far more likely to report that they had blocked a tool than employees were to report that they knew of a blocked tool, suggesting that most banned apps would never have been a problem anyway.

Second, when applications that employees wanted to use were blocked, they were just as likely to find a workaround for accessing those tools as they were to stop using them.

To that end, we asked employees who use unapproved tools why they do so. Given some of our internal debates over which AI models are the best, we expected it would come down to the quality or speed of answers. But the most common reason was much simpler: they used unapproved tools because it was easier. This finding is good news for security teams, as improving access to approved AI tools is something within their control.

AI Tool Use Despite Restrictions

Only about 10% of employees reported AI applications being blocked for them. Of those encountering blocked app, almost half reported workarounds to keep using the blocked tools.

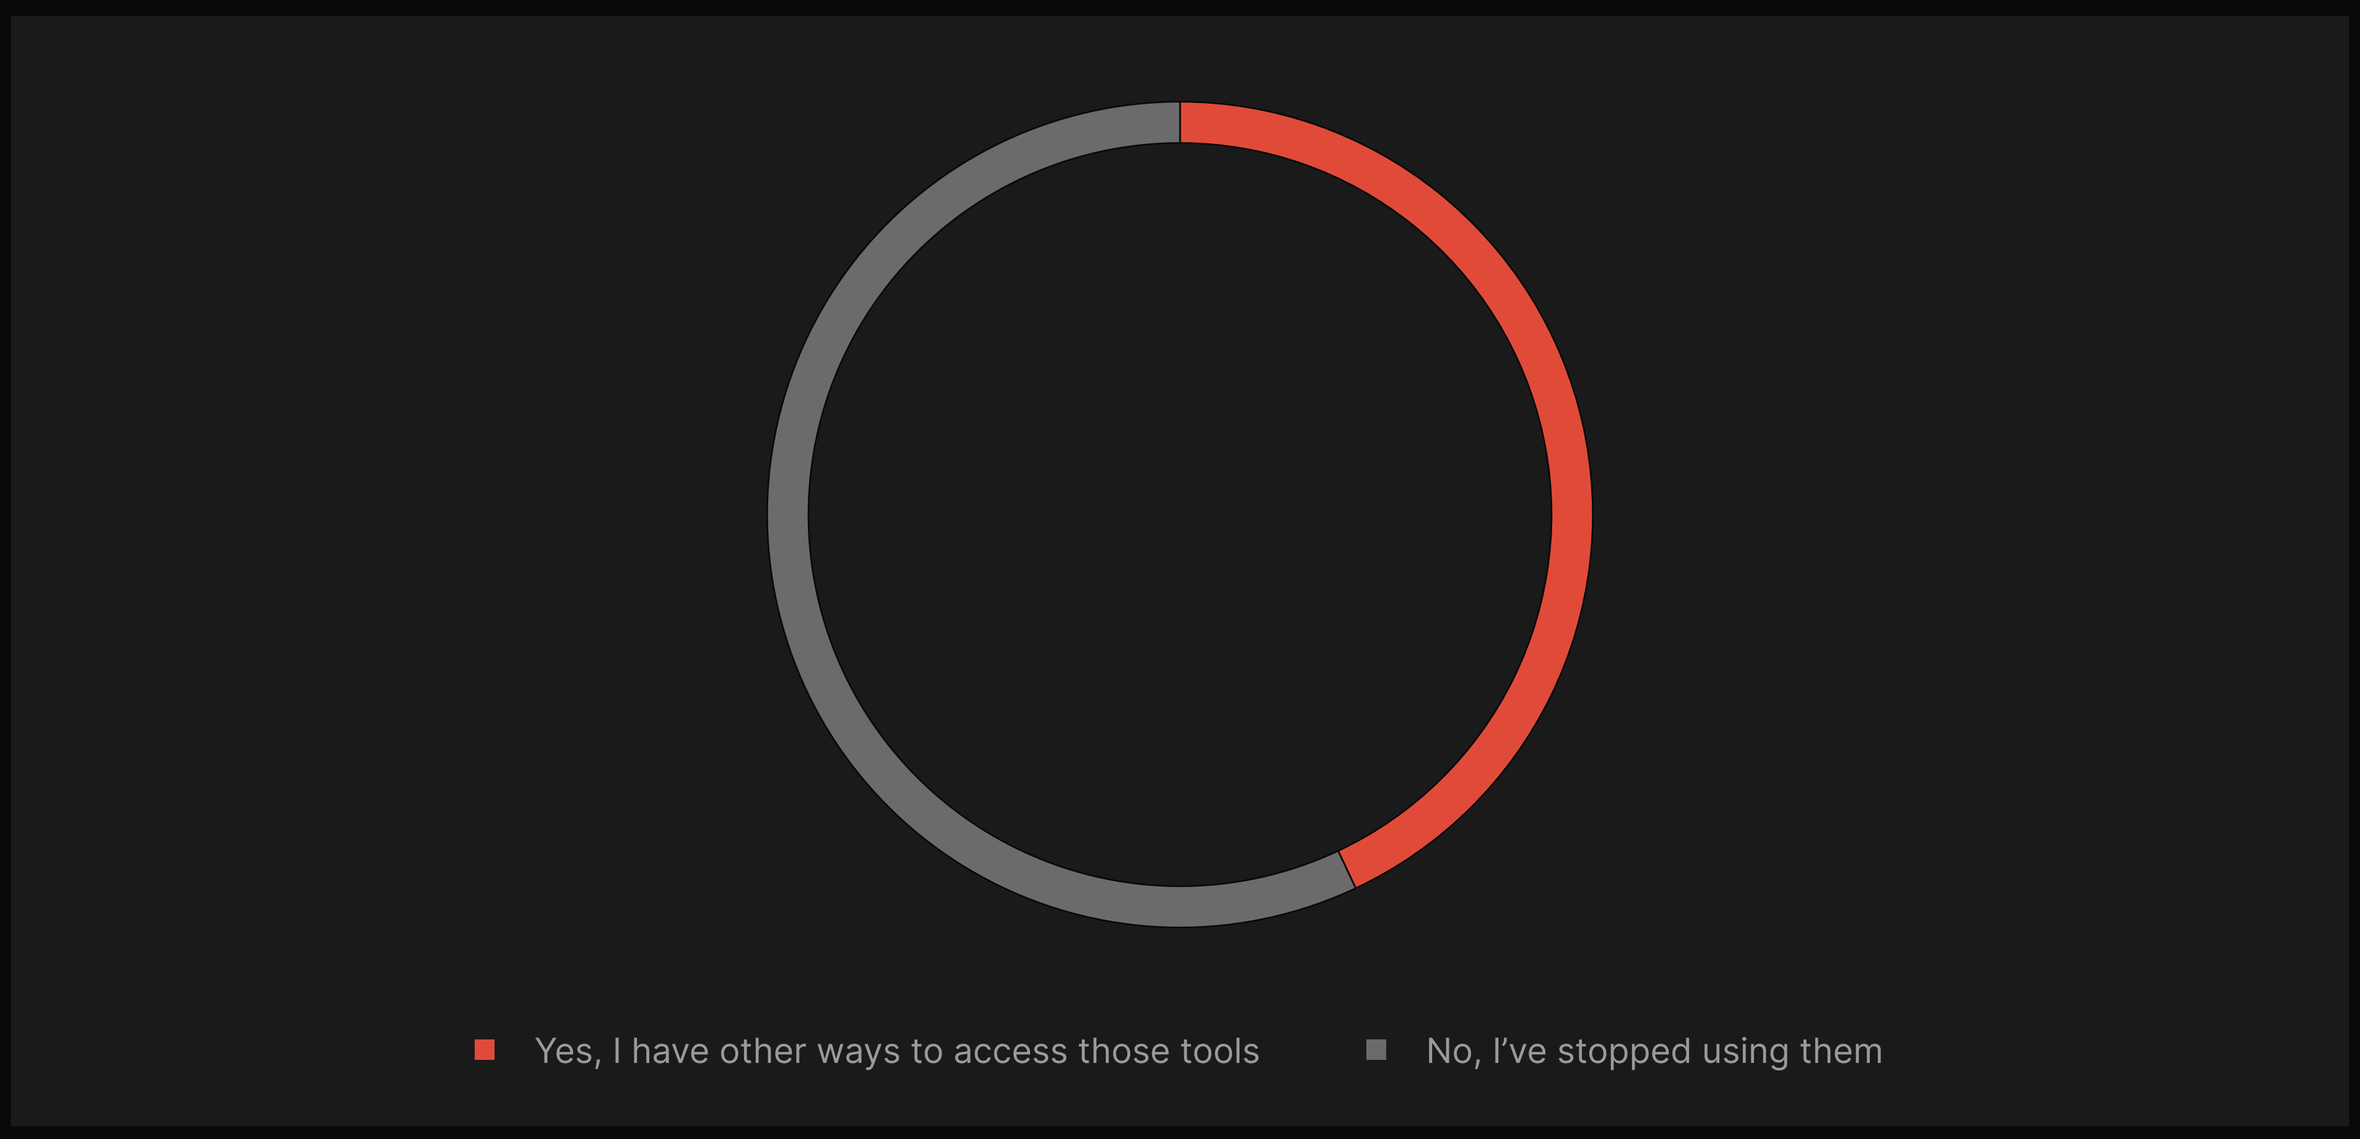

Security Orgs Blocking AI Tools by Region

About half of security organizations around the world report blocking some AI applications, ranging from 39% of Canadian to to 66% of American respondents.

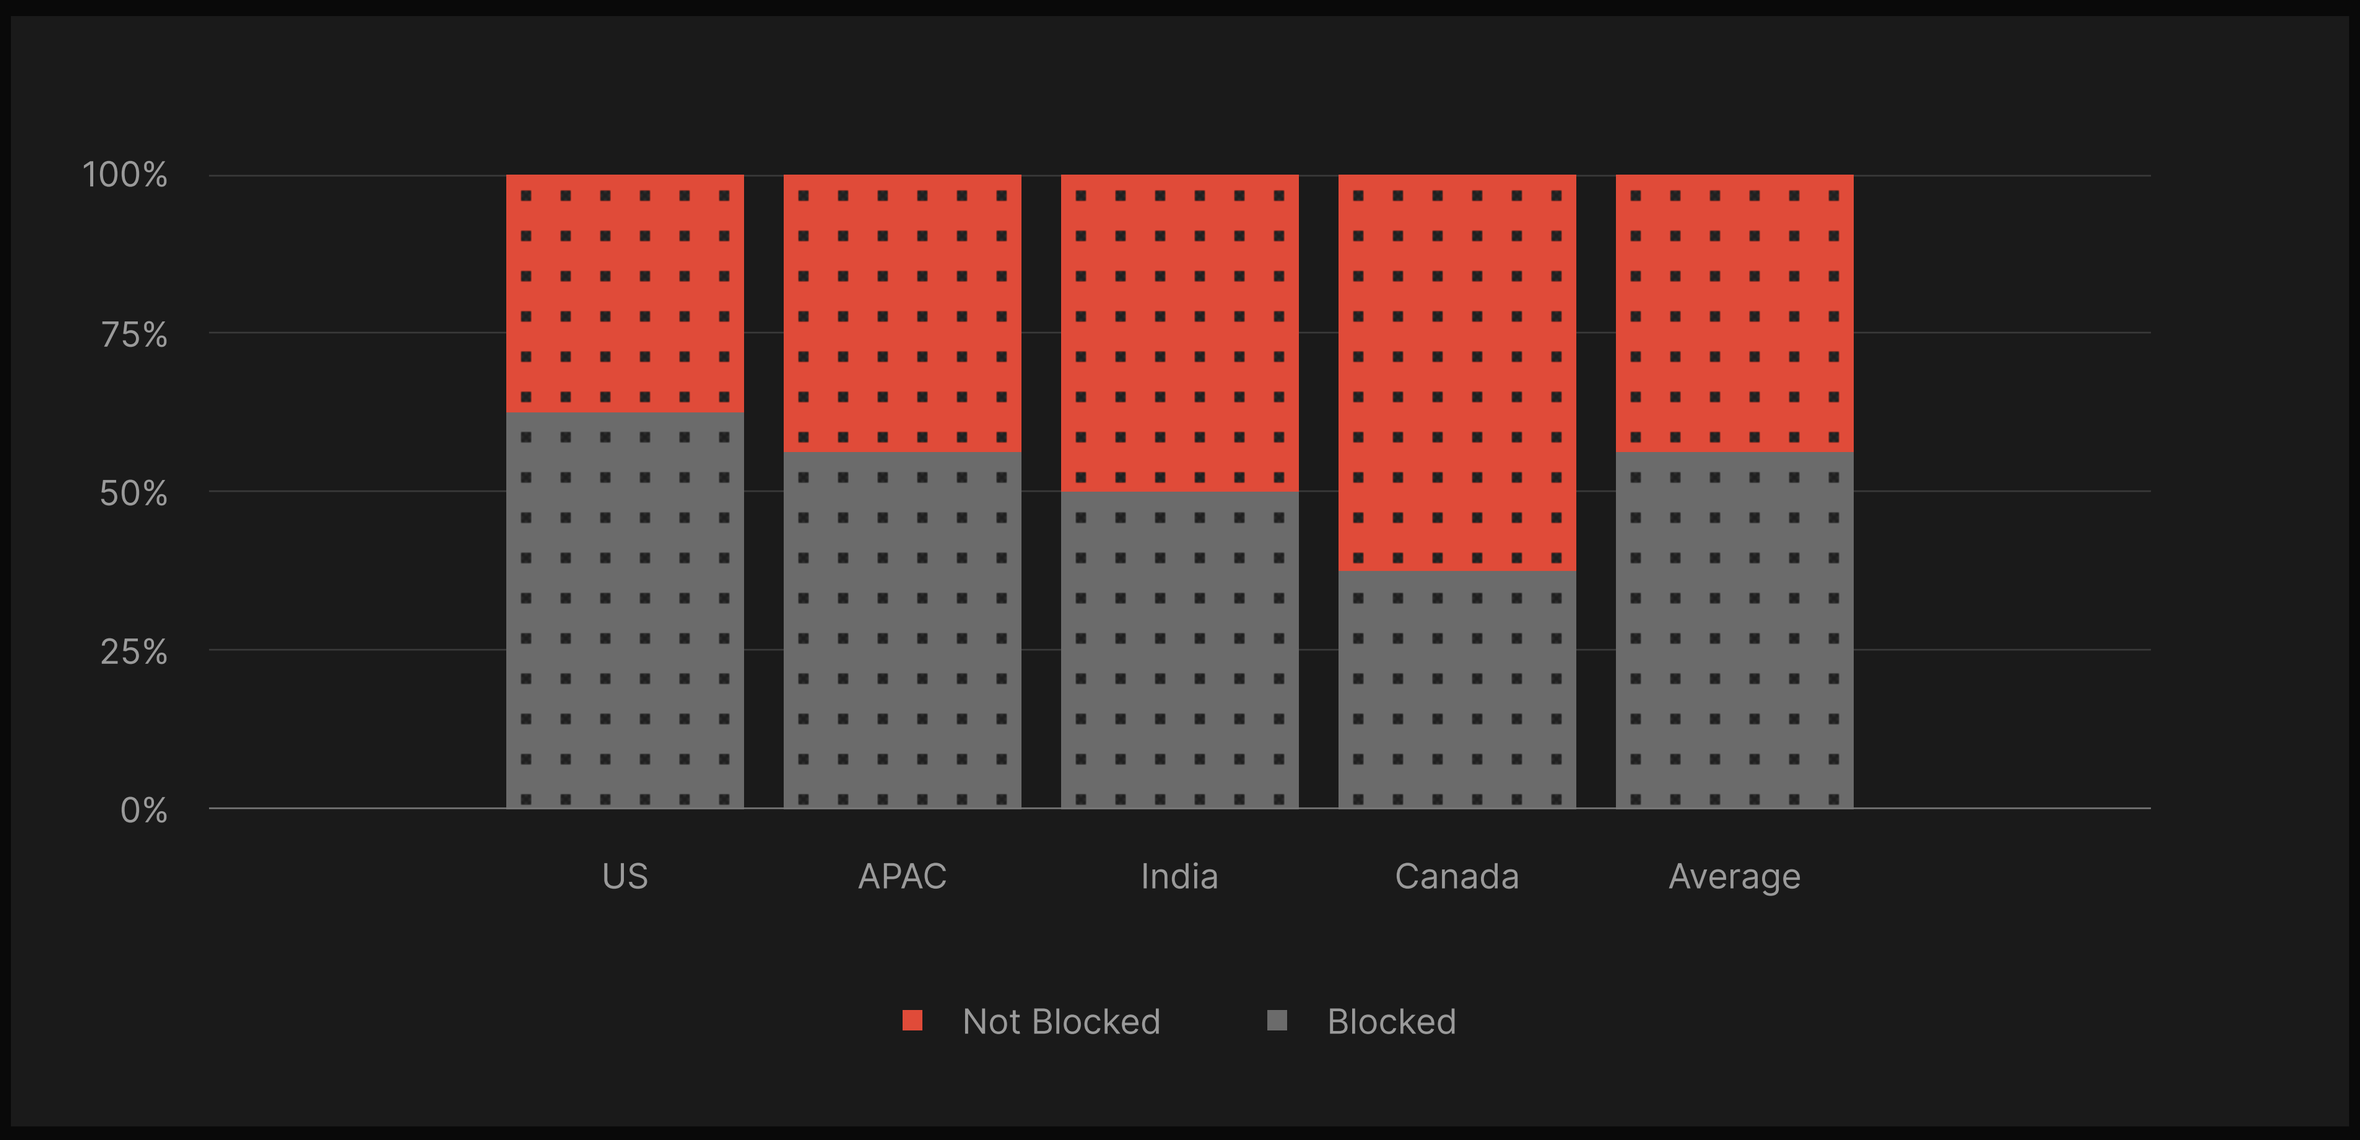

Awareness of Blocked AI Tools

On average, only 10% of workers reported AI applications being blocked in their organization, compared to the 56% of security leaders who said they blocked AI tools. Application blocking appears to mostly prevent imagined behavior, not real shadow AI.

Reason for Workforce Shadow AI Use

Employees don’t choose unapproved tools because they’re better or faster, they choose them because they’re easier. Security teams can beat shadow AI by competing on convenience.

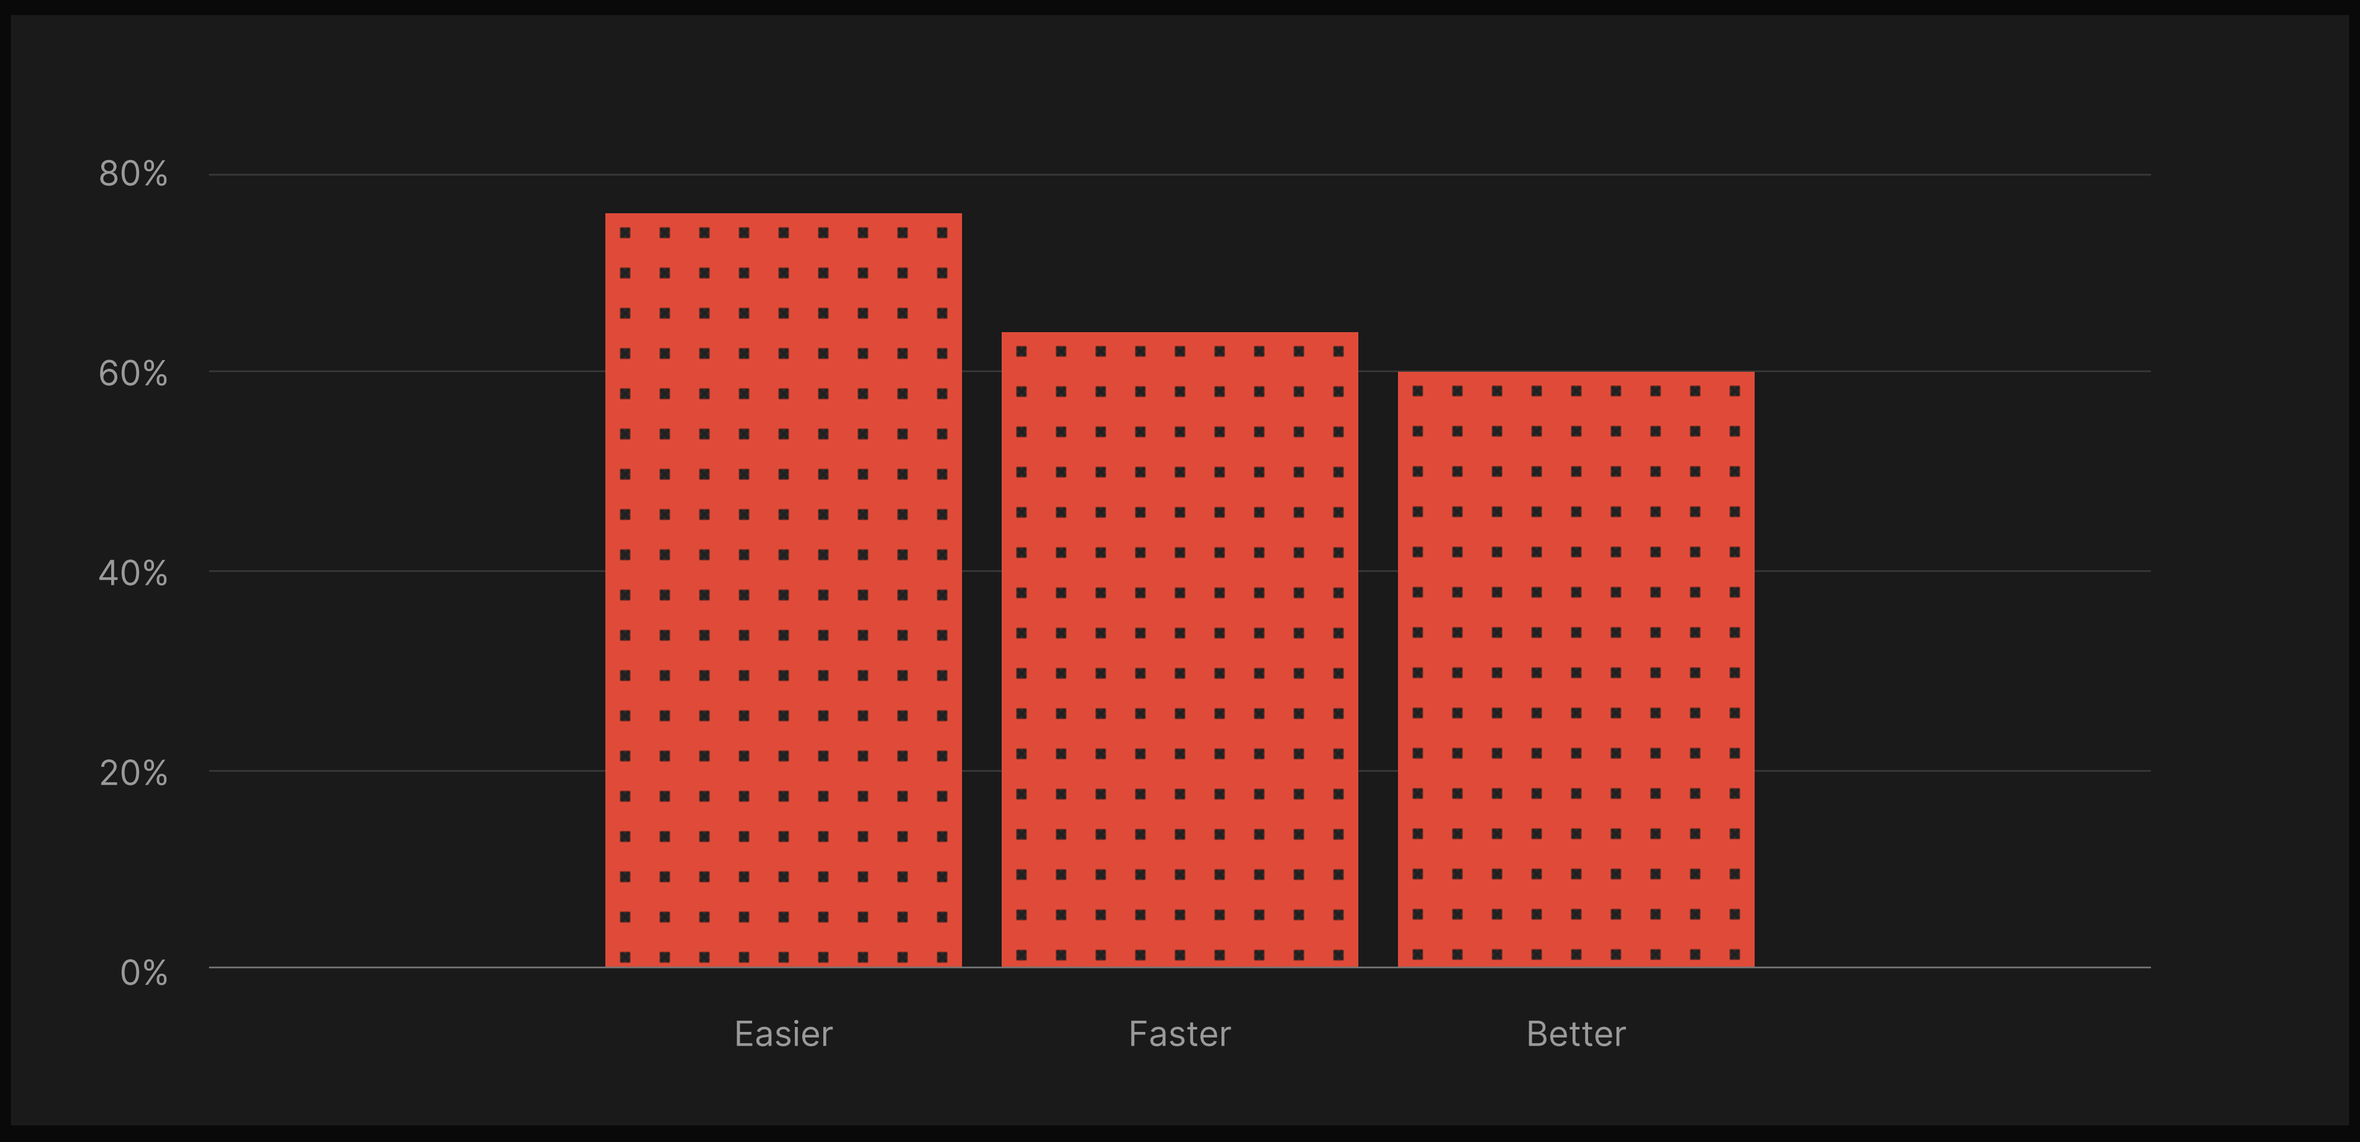

Building the reactor without the meltdown.

The embedding of AI into our daily lives may feel like an overnight eruption, but that sudden cascade was only made possible by the slow tectonic build of scientific research. Now we have a new research responsibility: developing the blueprint to harness this power productively. How do we build the nuclear reactor and avoid the meltdown?

The value offered by AI is immediate. If anything, it is too fast. The time needed to test, validate, and understand AI responses feels wasteful—maybe even primitive in contrast to the execution time of an advanced model. However, thanks to the efforts of researchers studying AI in the workplace, we know that without taking the time to plan properly, AI implementation projects are likely to fail. We also know where organizations need to invest time to ensure projects succeed. It may be less fun than vibe coding a whole business plan, but there is a substantial financial incentive to listen to the research on AI in the workplace.

Information security is in a similar situation, as the potential for productivity grates against the needs of data governance. Again, the necessary solution is to study the problem. Fact: Most employees are using unapproved AI tools. For many organizations, this discovery is where the investigation ends, usually resulting in reactive bans. Though this realization should only be the beginning of our inquiry, not the end. By exploring the motivations of the workforce, we can find safe methods to channel their curiosity for the net benefit of the business.

In this survey, we revealed the profile of the AI power user: governance-aware but defiant. They know the rules, but choose to use their own tools anyway. With that framing of the problem, we can stop pushing counterproductive “solutions” that drive users further into the shadows. Instead, the best solution is to prioritize visibility to maintain awareness of user activity.

Achieving this level of visibility requires a two-pronged strategy: one part cultural, one part technological. First, we must stop treating employee curiosity as a threat and provide safe, approved channels and training for continued exploration. By encouraging individuals to experiment and discover how AI can genuinely impact the business, we transform shadow AI from a liability into a governed engine for innovation.

However, this cultural shift can only be safeguarded by a simultaneous investment in modern human risk management. To encourage exploration without sacrificing governance, security teams must have a way to monitor AI usage and identify new risks in real time.

This visibility provides the crucial awareness needed to guide employees, apply smart policies, and protect data, all without resorting to the hard-block tactics that drove them into the shadows in the first place.

UpGuard

See it. Guide it. Secure it.

This study highlights a growing gap between AI use and organizational control. As shadow AI expands, so does human risk. UpGuard User Risk bridges that gap—providing continuous visibility into workforce behavior, identifying risk at its source, and enabling safe, compliant AI adoption.

See how User Risk works →